DIY Saturn-Like Hexagonal Jet: A Rotating-Tank Annular Shear Experiment for Polygonal Rossby Waves

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 20, 2026

Saturn's north-pole hexagon is one of the most visually arresting structures in planetary science: a persistent, nearly perfect six-sided boundary wrapped around a polar vortex. While the real phenomenon lives in a deep, rapidly rotating atmosphere, its core ingredients can be mimicked on a tabletop. This paper shows how to build a DIY laboratory analogue that produces a Saturn-like annular jet and then coaxes that jet into a stable six-lobed (m = 6) polygonal meander. The central idea is simple: place a shallow layer of fluid in a rotating tank to enforce near–two-dimensional, geostrophic flow, then drive a narrow shear layer at a chosen radius using a differentially rotating ring. As the shear strengthens, the circular jet destabilizes into a standing wave that can lock into a polygonal pattern, including a hexagon, over a surprisingly wide range of operating conditions. We present an end-to-end build plan, from mechanical layout and safe low-voltage power to dye visualization, spin-up protocol, and parameter tuning with dimensionless controls (Rossby and Ekman numbers). The goal is not just a striking demo, but a repeatable experiment that helps readers feel how planetary-scale order can arise from simple rotational physics.

Keywords: Saturn hexagon, Saturn north pole, polar jet, annular jet, rotating tank experiment, geophysical fluid dynamics, Rossby waves, barotropic instability, shallow water, differential rotation, Ekman number, Rossby number, DIY laboratory analogue, polygonal vortex, flow visualization, dye tracing, parameter tuning, planetary atmospheres, persistent patterns.

A collaboration with GPT-5.2-Thinking. 73 pages. CC4.0.

Outline

1. The Saturn Hexagon as a Laboratory Target
1.1 What “Saturn-like” means in a tabletop experiment
1.2 The difference between a polygonal core vortex and a polygonal annular jet
1.3 Why persistence and phase-locking matter more than perfect geometry

2. Physical Intuition: Why Jets Become Polygons
2.1 Rapid rotation and the emergence of quasi-2D flow
2.2 Shear layers, wave selection, and why a small integer wins
2.3 From instability to equilibrated polygon: the “locked wave” picture

3. The Minimal System That Produces an Annular Jet Hexagon
3.1 Rotating frame: tank, platform, and co-rotating lid
3.2 Differential driver: the surface ring that creates the jet at radius R
3.3 Why a narrow jet is essential for clean polygon formation

4. Design Parameters and Scaling Laws You Actually Use
4.1 Geometry choices: tank radius, jet radius, ring width, fluid depth
4.2 Rossby number as the forcing knob
4.3 Ekman number as the “smoothness versus noise” knob
4.4 Practical scaling: how to downsize without losing the hexagon

5. Bill of Materials and Tools
5.1 Tank and rotating platform options (budget to robust)
5.2 Ring assembly options: bearing, seal, and drive coupling
5.3 Motors, controllers, power, and wiring on a rotating frame
5.4 Low-cost measurement tools: tachometer, strobe, smartphone video

6. Mechanical Build: A Repeatable Annular Ring Drive
6.1 Building the co-rotating lid to suppress surface slosh
6.2 Mounting a concentric ring with low wobble and low friction
6.3 Drive strategies: belt, friction wheel, gear, or magnetic coupling
6.4 Alignment and leveling: how small errors kill symmetry

7. Electrical and Safety Architecture
7.1 Why low-voltage DC is the right choice
7.2 Power on a rotating platform: battery versus slip ring
7.3 Splash protection, guarding, and safe operating practices
7.4 A pre-run checklist to prevent the common failures

8. Fluid Preparation and Visualization
8.1 Water versus water–glycerin mixes: choosing viscosity
8.2 Dye injection methods that reveal the annular jet cleanly
8.3 Tracers and lighting: maximizing contrast for top-down imaging
8.4 Optional upgrades: UV fluorescein, polarizers, simple PIV workflows

9. Operating Protocol: From Spin-Up to Hexagon
9.1 Spin-up to solid-body rotation: how long and how to tell
9.2 Turning on differential rotation without triggering turbulence
9.3 Establishing the jet at radius R before chasing polygons
9.4 The tuning sequence that biases toward m = 6

10. Parameter Tuning for m = 6
10.1 A practical dialing guide: Ω first, then ω
10.2 Reading the pattern: how to count lobes reliably
10.3 If you get 7–8 lobes: how to move toward 6
10.4 If you get 5 or fewer: how to move upward to 6
10.5 Noise control: when to change depth, viscosity, or lid clearance

11. Diagnosing “Not Saturn-Like Yet” Outcomes
11.1 Polygonal core vortex: why it happened and how to avoid it
11.2 Axisymmetric shear with no polygon: likely causes and fixes
11.3 Turbulent breakdown: how to back out without losing the run
11.4 Wobble and drift: mechanical sources versus fluid sources

12. Data Collection and What Counts as Success
12.1 Documenting stability: duration, rotation rate, and phase-locking
12.2 Quantifying the jet: radius, width, and wave amplitude
12.3 Simple uncertainty handling for DIY instrumentation
12.4 A “minimum publishable dataset” for reproducible DIY science

13. Interpretation: What This Teaches About Planetary Atmospheres
13.1 What the analogue captures well (and what it cannot)
13.2 Persistence, forcing, and the meaning of a selected wavenumber
13.3 Why a north-pole hexagon can be stable without being rigid
13.4 How the lab setup reframes the question “why a hexagon?”

14. Extensions and Variants
14.1 Two-ring systems and more Saturn-like shear profiles
14.2 Bottom topography or thermal forcing as additional realism

14.3 Scaling up: when a larger tank dramatically improves clarity
14.4 Classroom and outreach versions that still preserve the physics

15. Conclusion: Building a Persistent Hexagon as Hands-On GFD
15.1 One-paragraph synthesis
15.2 Practical takeaways: the shortest path to a stable m = 6
15.3 What to explore next: controlled experiments you can run at home

16. References
16.1 Planetary observations and Saturn hexagon overview sources
16.2 Rotating-tank annular jet and polygonal wave experiments
16.3 Shallow-water, Rossby wave, and barotropic instability references
16.4 DIY build references: motors, slip rings, safety, measurement tools

Art

1. The Saturn Hexagon as a Laboratory Target

The fascination of Saturn's north-pole hexagon is not that it is a geometric novelty, but that it is a durable, dynamical boundary. It looks like a clean polygon drawn on the planet, yet it is made of moving gas. That combination, crisp shape plus living fluid motion, is exactly why it makes such a good laboratory target. A good lab analogue is not a miniature Saturn. It is a deliberately simplified system that preserves the few ingredients required to generate the same kind of organized, persistent structure from rotation and shear.

In practice, a tabletop "Saturn-like hexagon" is best thought of as a particular class of pattern in a rotating fluid: a narrow, fast jet organized around a polar region, which destabilizes into a low-mode wave and then locks into a persistent polygonal meander. The lab version aims to reproduce this class of behavior, not the full planetary environment. On Saturn, the hexagon sits within a deep atmosphere, embedded in complex radiative forcing, seasonal chemistry, and three-dimensional stratification. In the lab, we remove almost all of that and ask a more surgical question: if you provide a rapidly rotating frame plus a well-defined annular shear layer, can the fluid self-organize into a stable, integer-lobed polygon in the jet? The answer is yes, and that "yes" is enough to justify a DIY build as a serious demonstration of geophysical fluid dynamics, not a gimmick.

The practical reason to frame the hexagon this way is that it prevents the experiment from collapsing into "shape-chasing." If the goal is simply to see six corners somewhere in water, there are dozens of easier pattern-forming tricks. But those tricks often produce hexagonal tilings or surface ripples, not a Saturn-like jet boundary. The laboratory target here is specifically a jet that becomes polygonal through dynamics that resemble planetary jet instabilities. That clarifies what to build, what to tune, and what counts as success.

1.1 What "Saturn-like" means in a tabletop experiment

"Saturn-like" in this context means the experiment should satisfy three requirements.

First, the flow should be strongly constrained by rotation. You want a regime where the fluid behaves "geophysically," meaning it tends to become columnar and nearly two-dimensional in the horizontal plane, with the dominant balance between pressure gradients and Coriolis effects in the rotating frame. In ordinary everyday water motion, inertia and viscosity dominate and you get messy, quickly decaying swirls. In a rapidly rotating tank, the dynamics become disciplined: large-scale flow organizes, and waves and jets persist in ways that look qualitatively like planetary atmospheres. The experiment does not need to match Saturn's absolute rotation rate, but it must be in the same dynamical category where rotation matters enough to shape the patterns.

Second, the pattern must live on an annular jet rather than on a free vortex core. Saturn's hexagon is a deformation of a polar jet stream, not simply the boundary of a bathtub vortex. The tabletop "Saturn-like" target is therefore a narrow shear layer at a chosen radius, with relatively different angular velocities on the inside and outside of that radius. A differentially rotating ring at the surface is a clean way to create this. It creates a controlled annular jet with a tunable shear profile. This jet is where the polygon should form.

Third, the polygon must be a coherent wave pattern that can persist and rotate as a structure, rather than being a transient splash artifact. In other words, the corners should not be momentary. They should stabilize into a repeating wave with a clear integer mode count. A Saturn-like analogue is a dynamical attractor: once established, it remains visible over many rotations and survives small perturbations.

In this definition, "Saturn-like" does not mean the polygon must be perfectly regular or stationary relative to the lab room. It means the polygon is a robust, selected mode of a rotating annular jet, held together by the system's balances, not by careful shaping of boundaries or by optical tricks. The experiment is successful when you can say, with a straight face, that the pattern is not drawn, not forced in its shape, and not accidental. It is selected by the fluid.

1.2 The difference between a polygonal core vortex and a polygonal annular jet

This distinction is crucial because it is the most common way a DIY build goes wrong while still looking "cool." Many rotating-fluid experiments produce polygons, but they produce them in different physical locations and for different reasons. If you don't distinguish these, you can accidentally build the wrong experiment and then wonder why it doesn't teach you what Saturn teaches.

A polygonal core vortex is the "spinning-dimple" phenomenon: a central depression forms on the free surface, and the boundary of that depression becomes polygonal at certain rotation rates. This is visually dramatic and relatively easy to produce with a single motor and a rotating disk. The polygon appears because the vortex core and its surrounding shear layer become unstable and develop a small number of lobes. It is an instability of a swirling flow around a central core.

A polygonal annular jet is different. Here the defining object is not the vortex core but a jet ring at some radius, typically created by differential rotation between an inner and outer region. The polygon appears as a meander of the jet itself, a standing wave around the ring. You can have a strong central vortex in the system and still have the Saturn-like polygon living out on the jet, but the essential point is that the polygon is tied to the jet's shear layer, not to the shape of a central free-surface depression.

Why does this matter? Because the lab analogue you want is the latter. Saturn's hexagon is most naturally understood as a wave pattern on a circumpolar jet. If your tabletop polygon is just the edge of a vortex dimple, you have built something closer to a bathtub vortex demonstration. It may share mathematical motifs with jet meanders, but it is not the same class of object. It won't teach the same lessons about annular shear, wave selection on jets, and stable equilibrated modes.

There are also practical differences in how you tune each system. Core-vortex polygons often appear at relatively higher swirl and can be sensitive to surface effects, container geometry, and free-surface deformation. Annular jet polygons depend more cleanly on the differential rotation profile and on suppressing unwanted surface gravity waves, because the jet is a subtle structure that can be washed out by sloshing. In a Saturn-like build, you go out of your way to damp surface waves with a co-rotating lid, and you try to control where the shear lives by using a ring driver that defines the jet radius. Those are design choices you only make if you care about annular jets.

A concise way to remember the difference is this: in a core vortex polygon, the polygon is the boundary of a hole; in an annular jet polygon, the polygon is the boundary of a stream. Saturn is a stream.

1.3 Why persistence and phase-locking matter more than perfect geometry

If you judge success by how "straight" the sides look, you will miss the entire point of the phenomenon and you will also make the experiment harder than it needs to be. Saturn's hexagon is not a rigid wireframe. It is a pattern in a moving atmosphere. In the lab, the corresponding object is a wave that maintains its identity over time, even as fluid parcels move through it. That is what persistence means: the pattern is not made of the same water, it is made of the same dynamics.

A DIY experiment also has inevitable imperfections: small wobble in the platform, slight tilt, minor asymmetry in the ring, temperature gradients, dust, and electrical noise in the motor controller. If the hexagon only appears when everything is perfect, it is not a robust Saturn-like structure. A real planetary pattern persists under constant perturbation. The right criterion is therefore not geometric perfection but dynamical stability.

Phase-locking is the practical signature of that stability. In a rotating annular jet, many wave modes are possible in principle. The system "chooses" one of them as an equilibrated state. When it does, the wave has a definite rotation rate relative to the rotating frame, and its corners maintain a consistent angular separation. The pattern becomes trackable. You can mark one corner and it comes back to nearly the same position after many cycles. That is a stronger

claim than “I saw six bumps once.” It is the difference between a transient instability and a stable nonlinear wave.

From the standpoint of interpretation, phase-locking is also what makes the lab analogue philosophically satisfying. Saturn’s hexagon feels uncanny because it suggests that nature prefers a small integer. But the deep lesson is subtler: a rotating shear flow can settle into a low-dimensional attractor. The integer is not imposed by a hexagonal boundary; it is selected by the dynamics. In a successful experiment, the hexagon appears as one possible locked state among neighboring integers. You may see 7 or 5 first, then by tuning parameters you move into a region where 6 is the favored mode. That whole story, selection and locking, is closer to Saturn than any perfectly straight edge.

Finally, focusing on persistence prevents a common misinterpretation: that the goal is “a hexagon,” rather than “a Saturn-like hexagon.” In the lab, if you can repeatedly obtain a stable polygonal jet and demonstrate controlled transitions between mode numbers, you have done something more valuable than producing a single photogenic frame. You have built an experiment you can interrogate. You can change depth, viscosity, differential rotation, or lid clearance and watch the system respond. That is the bridge from DIY spectacle to genuine geophysical fluid dynamics.

If the reader takes only one criterion from this section, it should be this: count corners only after the pattern persists long enough that you can stop touching the controls, step back, and watch the polygon continue to exist on its own. That is when your tabletop becomes, in the relevant sense, Saturn-like.

2. Physical Intuition: Why Jets Become Polygons

A polygonal jet is not a magic trick and it is not “geometry imposed on water.” It is what a rotating fluid does when you give it a narrow shear layer and just enough freedom to express an instability without dissolving into turbulence. The reason this matters for a DIY build is that it replaces trial-and-error with a mental model. Once you understand why the jet wants to become a small-integer wave, you start tuning the experiment in the right direction: you strengthen rotation to discipline the flow, you shape a narrow shear layer so it has a clean instability pathway, and you aim for parameters where the instability saturates into a coherent pattern instead of shredding into noise.

At a high level, the story has three phases. First, rapid rotation forces the flow into a geophysical regime where horizontal motions dominate and the fluid becomes reluctant to vary along the rotation axis. Second, the annular jet becomes a shear layer that supports waves, and among the many possible waves, a particular mode number is naturally selected because it fits the jet’s

width, curvature, and strength. Third, nonlinearity prevents the instability from growing forever: the wave grows until it reorganizes the jet into a new steady state, a polygonal meander that is self-maintaining and phase-locked.

That entire narrative is what “Saturn-like” means in physical terms. You are not building a hexagon. You are building a system that, when pushed past a threshold, prefers to settle into a low-mode wave on a jet.

2.1 Rapid rotation and the emergence of quasi-2D flow

When a fluid is not rotating, three-dimensional turbulence is the default fate of forced motion. Eddies cascade, break, interact, and quickly become complicated. You can still see large structures, but they are often short-lived and dominated by the details of stirring. In a rapidly rotating frame, that changes dramatically. Rotation imposes a kind of directional discipline: it makes it hard for the fluid to move in ways that require significant variation along the rotation axis.

The intuitive mechanism is simple. In the rotating frame, moving parcels feel an apparent sideways deflection proportional to the rotation rate. When rotation is strong compared to the characteristic timescale of the motion, the deflection acts quickly, and flows tend to organize into patterns where pressure gradients and this sideways deflection balance. In that regime, the fluid does not behave like “water in a bowl.” It behaves like a simplified atmosphere or ocean: large-scale, horizontally organized motions that can persist.

A helpful picture is to imagine the fluid trying to form a tall eddy column. In a non-rotating system, the column can break apart vertically and twist in three dimensions. In a rotating system, vertical coherence is favored. The flow becomes “columnar,” meaning that velocity patterns look similar at different depths. This is the essence of the quasi-2D behavior. The fluid still has depth and vertical structure, but the dominant dynamics can be interpreted in a horizontal plane with depth acting more like a parameter than a full third dimension.

For a DIY Saturn-like hexagon, this matters for two reasons. First, it makes the jet look like a real jet rather than a smeared swirl. A jet is a coherent band of flow with a sharp change in velocity across it. If the flow is too three-dimensional, the jet gets shredded. Second, it allows waves on the jet to exist long enough to be seen. In an undisciplined flow, instabilities grow and then immediately tumble into turbulence. In a rotating flow, the instability can grow into a coherent wave and then saturate into a stable pattern.

This is why “spin-up” is not a procedural nuisance but a physical necessity. You are giving the system time to align with the rotating frame, damp out transient sloshing, and enter the regime where geostrophic-like balance can dominate. The difference between a messy vortex and a

stable polygonal jet is often nothing more mysterious than “rotation strong enough, and patience long enough.”

2.2 Shear layers, wave selection, and why a small integer wins

Once you have a rotating frame, the next ingredient is a narrow shear layer. In the annular ring experiment, the shear layer is the jet itself. One side of the jet is moving faster than the other, and that velocity difference is concentrated across a relatively small radial width. That is the energy reservoir for the instability. The jet is stable as long as perturbations cannot extract energy efficiently from the shear. Once conditions cross a threshold, small disturbances can draw on the shear and grow.

But why does the jet grow a pattern with a small integer number of lobes rather than something complicated? The answer is wave selection. The jet is not just a band of flow; it is also a waveguide. It supports disturbances that propagate around it, and these disturbances have characteristic spatial scales. If a disturbance is too small-scale, viscosity and three-dimensional effects damp it. If it is too large-scale, it may not “fit” the curvature and structure of the jet in a way that extracts energy efficiently.

The annular geometry adds a second constraint: periodicity. A wave that travels around a ring must close on itself. That means the circumference of the jet must contain an integer number of wavelengths. That integer is exactly the polygon count you see. If six wavelengths fit around the ring, you see six lobes. If five fit better, you see five. The ring forces the system to choose a discrete set of possibilities.

This is why “small integer” modes are natural. Large integers correspond to short wavelengths, which are more easily damped and more sensitive to imperfections. Very small integers correspond to very long wavelengths, which may not couple well to the jet’s shear profile. In many parameter ranges, the “sweet spot” is a low mode number like 5, 6, 7, or 8. The sweet spot shifts as you change forcing strength, depth, viscosity, and jet width. This is also why the experiment often walks through different polygons as you tune it: you are moving through a landscape of mode preference.

There is also an important nonlinear twist. The mode that grows fastest at onset is not always the mode that remains stable at large amplitude. Near the threshold, the jet may prefer a higher mode number, because small-scale structures can grow more easily. As you increase forcing, the jet can reorganize, and the stable equilibrated mode number may decrease. This leads to a characteristic observation: as you crank up differential rotation, you may see the pattern “step down” from 8 to 7 to 6 to 5, rather than smoothly changing. That stepping is the discrete nature of mode selection in a ring.

In a DIY setting, this viewpoint is extremely useful because it tells you what to do when you get the “wrong” polygon. If you see 7 or 8 lobes, you are not failing. You are in the neighborhood of the correct regime. Small parameter nudges can shift the preferred integer. The experiment is not “produce a hexagon or nothing.” It is “navigate a discrete mode diagram until m = 6 is the stable attractor.”

2.3 From instability to equilibrated polygon: the “locked wave” picture

Instability is the beginning of the story, not the end. If the only thing happening were instability, the pattern would grow until it broke and became turbulence. The reason a stable polygon exists is that the instability saturates. The growing wave reaches an amplitude where it changes the base flow that fed it. The jet is not merely a passive stage; it is reshaped by the wave. Eventually, the system finds a configuration where the wave can persist without continuing to grow uncontrollably.

The “locked wave” picture is a way to describe that final configuration. Imagine the jet as a circular band. A wave develops on it, creating alternating bulges inward and outward. Those bulges are not just shape deformations; they are associated with vortical structures that ride along the jet. The wave pattern organizes the shear so that the energy extraction from the jet is balanced by dissipation and by the self-induced changes in the flow. At that balance point, the wave’s amplitude becomes steady. The pattern is now an equilibrated state.

In a rotating system, equilibrated waves can behave like rigid objects even though the fluid is moving. The wave has a coherent phase, and that phase can drift at a constant rate relative to the rotating frame. This is what phase-locking looks like in practice: the corners of the polygon rotate around the tank at a steady angular velocity, and the spacing between them remains nearly constant. The pattern is not stationary in the lab frame, but it is stable in the dynamical sense. It returns to itself.

This is also why the co-rotating lid and careful damping are so important. The locked wave is a subtle balance. Surface sloshing adds competing wave modes that can disrupt phase coherence. Mechanical wobble adds periodic forcing that can either entrain the pattern or destabilize it. If those perturbations are too strong, the system cannot settle into the locked state and instead keeps wandering among transient shapes. The goal of good experimental design is to make the instability’s nonlinear saturation the dominant pathway, not external vibration.

The locked wave picture also clarifies why “perfect geometry” is the wrong criterion. In a real nonlinear equilibrated wave, the lobes are not perfectly identical. There can be mild asymmetry, slight amplitude differences, or slow drift in where the wave crests sit relative to the ring. Yet the mode number remains stable and the pattern persists. That persistence indicates you have achieved the key outcome: a low-dimensional attractor in a rotating shear flow.

For a DIY builder, the practical use of this picture is that it tells you what to aim for when tuning. You want to see the wave emerge, grow, and then stop growing. You want the corners to become crisp and then remain crisp. You want the polygon to have a stable rotation rate, not a jittery wandering. When those conditions are met, you have succeeded in creating not just a "hexagon," but the specific dynamical phenomenon that makes Saturn's hexagon worth reproducing in the first place.

3. The Minimal System That Produces an Annular Jet Hexagon

A Saturn-like annular-jet hexagon is one of those experiments that looks impossibly exotic until you realize how few ingredients it truly requires. You need a fluid layer, a rotating frame strong enough to "geophysicize" the dynamics, and a way to impose a controlled shear at a chosen radius without introducing chaotic surface slosh. Everything else is refinement. If you can build these three elements cleanly, you can create a regime in which the flow naturally selects a low-integer wave mode and settles into a persistent polygonal meander of the jet.

This section describes the minimum system in functional terms, not as a shopping list. The key design idea is to split the experiment into two distinct motions that are easy to confuse: a global rotation that defines the reference frame, and a differential rotation that defines the jet. The global rotation sets the "planet." The differential rotation creates the "wind." Once those are separated and controlled, the rest of the experiment becomes a matter of tuning.

3.1 Rotating frame: tank, platform, and co-rotating lid

The rotating frame is not a convenience; it is the experiment's foundation. Without strong, steady background rotation, the fluid will not remain organized long enough for a polygonal jet to emerge and stabilize. There are three components: the tank, the platform, and the lid. Each does something physically necessary.

The tank is the container that defines the circular domain. It should be as round and symmetric as you can make it. The exact material is less important than its geometry and balance. A clear cylinder is ideal because you want top-down visibility and you want to diagnose the jet radius by dye. A tank diameter in the 30–60 cm range is a sweet spot for DIY. Smaller tanks can work, but the patterns become harder to see and mechanical imperfections become proportionally larger. A wider tank gives you more circumferential room, which improves mode clarity and makes the difference between "six-ish bumps" and a convincing hexagon.

The platform is the device that rotates the entire tank at a controllable angular speed Ω. What matters is not only the maximum speed but also stability and smoothness. A slightly uneven rotation rate is a periodic forcing signal, and periodic forcing can either entrain the wave or

destroy it. In practice, a pottery wheel or a robust turntable with variable speed control tends to be the most convenient DIY option because it provides torque and smooth rotation, but any platform can work if it rotates steadily, can carry the load, and does not wobble.

Leveling is not optional. A rotating tank that is even slightly tilted creates a preferred direction in the fluid and will bias the jet into drifting or precessing. The experiment relies on rotational symmetry so that the only preferred features are the ones generated by the dynamics themselves. If the tank is not level, you will get a “polygon” that is partly the fluid responding to gravity’s slight lateral component rather than to the intended instability.

The co-rotating lid is the component that most DIY builds omit, and it is often the difference between success and frustration. In a free-surface tank, the surface readily supports gravity waves and sloshing. Those waves inject competing modes into the flow, mask the subtle jet meander, and can trigger premature turbulence. A lid that rotates with the tank damps those waves by constraining the free surface and reducing the air–water interface’s ability to oscillate. The lid does not need to seal the tank like a pressure vessel, but it must be close enough to the surface to suppress slosh without contacting the fluid.

A good lid also stabilizes the experiment mechanically. If you can attach the lid rigidly to the rotating frame, it becomes a reference plane and can carry the differential ring drive cleanly. Conceptually, the lid turns your open bowl into a controlled rotating apparatus. You are trying to create something more like a planetary “cap” than a kitchen whirlpool.

In summary, the rotating frame must provide: steady Ω, good symmetry, minimal wobble, strong damping of surface waves, and a reliable reference plane for mounting the differential driver. If you achieve those, the flow stops behaving like stirred water and starts behaving like the simplified atmosphere you need.

3.2 Differential driver: the surface ring that creates the jet at radius R

Once the rotating frame is established, you still do not have a jet. You have a body of fluid tending toward solid-body rotation. The Saturn-like structure requires a narrow annular shear layer: a location where the angular velocity changes sharply with radius. That is exactly what the differential driver produces.

The differential driver is a ring at the surface (or very near the surface) that rotates at a small differential rate ω relative to the background rotation Ω. The ring’s job is not to “stir the water” in a general sense, but to impose a controlled mismatch between fluid inside and outside the ring radius. If the ring is coupled well to the fluid, it drags fluid near its radius slightly faster or slower than the surrounding region. The result is a jet-like shear layer centered around radius R, where R is the ring’s mean radius.

There are two key features of this approach.

First, it isolates the forcing to a chosen radius. If you rotate the bottom disk or the entire tank without differential forcing, the shear tends to be distributed, and the dominant structure may become a central vortex. With a ring, you can decide where the strongest shear lives. You can "place" the jet. That placement is essential because it gives the jet a well-defined circumference and therefore a well-defined set of admissible mode numbers.

Second, the ring produces differential rotation without requiring a second large rotating container. A common misconception is that you need an inner rotating cylinder and an outer rotating cylinder like a professional Taylor–Couette apparatus. That is one way, but it is not the minimum DIY path. A carefully mounted surface ring, rotating differentially, can produce a clean annular jet in a simpler build. This is the "minimal" trick.

For the ring to work, it must rotate smoothly, remain concentric, and couple to the water in a consistent way. Practically, that means low wobble, low friction bearings, and a ring edge geometry that interacts with the fluid. In many builds, the ring is a thin annulus that dips slightly into the water or sits just at the surface, creating shear by viscous coupling. The ring does not need to push water like a paddle wheel; it needs to establish a velocity gradient across a narrow radial band. Too aggressive a mechanical coupling can inject turbulence; too weak a coupling fails to produce a clear jet. The design sweet spot is a ring that shows discipline: it transmits differential rotation without acting like a blender.

The ring's differential rate ω is usually small compared to Ω, and that fact is important. The background rotation is doing the heavy lifting of organizing the flow; the ring is a perturbation that sets the jet. In successful Saturn-like regimes, the differential forcing is strong enough to destabilize the jet but not so strong that the entire flow becomes a chaotic shear-driven mix.

This separation of roles—Ω as the geophysical frame, ω as the jet forcing—also makes tuning intelligible. You can think of Ω as setting the flow's "stiffness" and ω as setting the jet's "stress." The polygon emerges when stress is high enough to excite a wave but stiffness is high enough to keep the wave coherent.

3.3 Why a narrow jet is essential for clean polygon formation

If the rotating frame is the stage and the ring is the actor, the narrow jet is the script. Without a narrow jet, the system may still become unstable, but it will rarely produce a clean, persistent polygon that you can count reliably. Instead you get broad meanders, multiple interacting bands, or turbulent breakup.

There are three reasons narrowness matters.

The first is wave clarity. A narrow jet behaves like a well-defined waveguide. Disturbances are concentrated on a band where the velocity gradient is strongest. That concentration makes the mode structure clearer: the jet develops a small number of large lobes rather than a fuzzy superposition of many small ones. A wide shear zone supports a larger variety of wavelengths simultaneously, which encourages mode competition and drifting patterns. If you want a crisp integer mode, you want the instability to "choose" one dominant scale. Narrow jets encourage that.

The second reason is mode selection by circumference. A polygon is fundamentally a circumferential wave count. The more sharply the jet is localized at radius R, the more meaningful that circumference becomes. In other words, if the jet has a clear radius, then "six lobes around the ring" is a well-posed statement. If the shear is spread across a wide radial region, then there is no single circumference; different parts of the flow prefer different effective radii and therefore different mode numbers. The result is a smeared, ambiguous polygon.

The third reason is saturation into a stable pattern. The equilibrated polygon is the nonlinear endpoint of an instability. For saturation to be clean, the instability must reorganize a specific structure—the jet—rather than rearranging the entire tank's flow. A narrow jet allows the instability to reshape the jet into a meandering path while leaving the broader interior and exterior regions relatively quiet. This separation supports persistence. If the shear is broad, the instability has too much freedom to couple into large-scale mixing and turbulence, which prevents the system from settling into a locked wave state.

In DIY practice, "narrow jet" translates into a few build choices.

It means the ring width should not be too large, and the ring should couple at a specific radius rather than over a broad area. It means the fluid depth should be chosen so that vertical mixing does not smear the shear layer. It means you want the lid close enough to suppress surface waves that would otherwise broaden the effective forcing. It means smooth rotation and careful alignment so that the jet is not constantly being thickened by wobble-induced stirring.

When you get this right, you will see a sharp dye band that traces the jet radius. That dye band becomes the visual carrier of the polygon. It will begin as a nearly circular ring, then develop gentle undulations, then sharpen into a small integer of lobes, and finally stabilize into a persistent meander. That progression is the signature of a narrow jet doing what it is supposed to do: offering a clean channel for a wave to grow and then lock in.

This is the minimal system's logic. You do not need a sophisticated laboratory to create a Saturn-like annular jet hexagon. You need a disciplined rotating frame, a controlled differential ring

forcing at a chosen radius, and a narrow shear layer that lets the instability express itself as a coherent low-mode wave rather than as general turbulence. Once those are in place, you are no longer guessing. You are tuning a system that is predisposed to produce the phenomenon you want.

4. Design Parameters and Scaling Laws You Actually Use

A DIY Saturn-like hexagonal jet is won or lost in the design parameters, not in philosophical intent. The most productive way to think about "design" is to separate what you can choose freely from what you cannot. You can choose geometry, depth, and where the jet lives. You can choose how fast the platform rotates and how hard the ring drives the shear. You cannot choose away viscosity, mechanical imperfections, or the basic demand that rotation must dominate enough to create geophysical behavior. Scaling laws exist to help you navigate those constraints with the fewest wasted runs.

This section is intentionally practical. The goal is not to derive fluid dynamics from first principles. The goal is to tell you which numbers matter, what they physically mean, and how to pick them so you land in a regime where an annular jet can destabilize into a low-mode polygon and then remain coherent long enough to be convincing.

Two dimensionless controls do most of the heavy lifting: the Rossby number and the Ekman number. The Rossby number tells you how hard you are forcing the jet relative to the disciplining effect of rotation. The Ekman number tells you how much viscosity and boundary friction will smear the flow and inject noise. Geometry sets the circumference, the curvature, and the available mode numbers. Together, these determine whether you get a stable hexagon, a different polygon, or turbulence.

4.1 Geometry choices: tank radius, jet radius, ring width, fluid depth

Geometry is where you decide what "six around the circle" will even mean. You are defining the waveguide and the set of allowed circumferential integers. The core choices are tank radius, jet radius, ring width, and fluid depth.

Tank radius sets the physical scale and how forgiving the experiment is. Larger tanks make everything easier: the hexagon is larger and clearer, mechanical wobble becomes proportionally smaller, and the circumference is long enough to support clean low-mode waves. Smaller tanks compress the pattern and demand higher precision. If your goal is a visually persuasive hexagon rather than a minimal proof-of-concept, prefer the largest diameter you can rotate smoothly and safely.

Jet radius is the most meaningful geometric parameter for the polygon. The polygon's mode number is literally "how many lobes fit around the jet circumference." That circumference is $2\pi R$, where R is the jet radius. A larger R increases the circumference and makes it easier for the jet to host a clean wave without the lobes becoming too tight and damped. In many practical builds, placing the jet at roughly half the tank radius is a good starting point. Too close to the center, and the dynamics begin to resemble a core vortex. Too close to the wall, and boundary effects and friction complicate the flow.

Ring width is how sharply you impose shear. The ring is your forcing instrument. A wide ring tends to create a broad shear zone rather than a narrow jet. A narrow ring concentrates the forcing and produces a sharper jet, which is exactly what you want for clean mode selection and stable polygon formation. In practice, you are choosing a width that is mechanically stable but still narrow enough to act like a localized driver. A useful mental ratio is ring width as a small fraction of jet radius. If your ring width becomes comparable to the radius itself, you are no longer driving an annular jet; you are driving an entire region.

Fluid depth is the fourth geometric choice, but it acts like a dynamical control. Depth sets the vertical scale of the flow, affects how strongly bottom friction matters, and influences how quasi-two-dimensional the behavior will be. Very shallow layers make the flow more sensitive to boundary effects and may damp the jet. Very deep layers can permit more three-dimensional motion and may require stronger rotation to keep the flow disciplined. A moderate depth, shallow enough to keep the flow organized but deep enough to avoid immediate viscous domination, tends to work best.

Depth also interacts with visibility. A deeper layer makes dye structures harder to interpret from above because you may be seeing vertical integration of motion rather than a surface-traceable jet. The lid and ring arrangement often encourages a surface-defined jet, so you want depth that supports that without creating excessive vertical mixing.

These geometric choices jointly define two practical outcomes: how sharp the jet looks when dyed, and how unambiguous the polygon count is when the wave forms. The most reliable visual cue that your geometry is working is a dye band that remains narrow and coherent at radius R for long periods, before it ever becomes polygonal.

4.2 Rossby number as the forcing knob

In the DIY annular-jet experiment, the Rossby number is the knob you turn when you want the jet to change its mode number. It is the dimensionless measure of how strong the differential forcing is relative to the constraining influence of background rotation.

At the level you actually use, Rossby number answers a simple question: are you gently perturbing a disciplined rotating flow, or are you slamming it hard enough that inertia takes

over? If the Rossby number is too small, the jet remains stable and nearly circular. If it is too large, the flow breaks into turbulence or becomes dominated by non-geostrophic motion that does not lock into a clean polygon. In the middle is where polygonal equilibria live.

The reason Rossby number is so useful is that it combines three things you directly control: jet radius, differential rotation speed, and background rotation. Increasing the differential ring speed increases the forcing. Increasing the background rotation increases the system's tendency to remain geostrophic and coherent. Changing depth changes the effective "stiffness" of the flow and the relative influence of forcing. Rossby number packages those into one scale-free indicator.

In practice, you use Rossby number as follows.

You choose a background rotation rate that produces a stable, quiet rotating frame. You then increase differential ring speed slowly until the jet begins to show a wave. That onset marks crossing a critical forcing threshold. You then continue nudging differential speed in small steps until the wave settles into a stable mode. If you get seven or eight lobes and want six, you adjust Rossby number slightly. If you get five, you adjust in the opposite direction. You do not need to compute Rossby number in the middle of the run, but thinking in Rossby-number terms keeps you from making the most common error: changing several things at once. Rossby number tells you that differential speed is the primary forcing knob; change that first before altering geometry.

Rossby number also explains why the experiment can feel discrete. You are not sliding continuously from one shape to the next. You are moving between stable integer modes. Within a given parameter band, one mode is an attractor. When you push forcing past a boundary, the attractor changes, and the system "snaps" to a different integer.

So, Rossby number is the forcing knob in the deepest sense: it is what moves you across the mode boundaries in the regime diagram, and it is what determines whether the instability saturates cleanly or collapses into turbulence.

4.3 Ekman number as the "smoothness versus noise" knob

If Rossby number is "how hard you push," Ekman number is "how cleanly the system can respond." Ekman number quantifies the role of viscosity and boundary friction relative to rotation. In the annular jet experiment, it acts like a smoothness-versus-noise knob because it controls how much the flow is smeared by frictional effects and how much unwanted secondary circulation is generated near boundaries.

In a rotating fluid, viscous friction near solid boundaries creates thin layers where the flow adjusts to satisfy no-slip conditions. Those layers feed secondary circulations that can either

stabilize or destabilize the jet and can inject noise that prevents clean phase-locking. When Ekman number is too large, viscous effects are not confined to thin layers; they reach into the bulk and smear out the jet, reducing the sharpness of the shear layer. When Ekman number is small enough, viscous effects are mostly confined to boundary layers and the bulk remains clean and organized.

For DIY design, Ekman number is influenced by three things: viscosity (which depends on fluid choice and temperature), background rotation rate, and fluid depth. Increasing background rotation reduces Ekman number. Increasing depth reduces Ekman number. Increasing viscosity increases Ekman number.

This gives you a practical interpretation: Ekman number tells you how much your experiment will behave like a disciplined rotating flow versus a viscously damped swirl. If the jet never becomes crisp and the dye always diffuses into a broad smear, you may be in a regime where Ekman effects are too strong. If the flow is overly lively and breaks into small-scale turbulence easily, you may be too far in the opposite direction, where viscosity is too weak to help the instability saturate cleanly. That is why viscosity adjustments, such as adding a modest fraction of glycerin, can sometimes improve the experiment even though it feels counterintuitive. You are not trying to eliminate viscosity. You are trying to place it in the role of gentle damper rather than dominant smearing force.

The “noise” in this phrase is not just thermal noise or measurement noise. It is the collection of unwanted dynamics introduced by imperfect boundaries, wobble, surface ripples, and small turbulent bursts. A lower Ekman number tends to reduce the extent to which boundary friction contaminates the bulk flow, which improves pattern coherence. But if you drive too hard at very low damping, you can also end up with a pattern that is unstable because nothing dissipates small disturbances. Thus, Ekman number is not simply “lower is better.” It is “low enough that the bulk behaves geophysically, but not so low that the system becomes hypersensitive.”

For a DIY builder, the practical rule is: if you can rotate faster without wobble, do that before you start changing fluids. Rotation rate is the cleanest way to reduce effective Ekman number. If rotation rate is limited, then depth is your next lever. Only then do you reach for viscosity tweaks.

4.4 Practical scaling: how to downsize without losing the hexagon

Downsizing is where many DIY attempts fail for reasons that feel mysterious until you interpret them through these parameters. The usual approach is to shrink the tank diameter because it is cheaper and easier to rotate. Then everything “should” still work because the physics is the same, right? The physics is the same, but the parameter regime is not automatically preserved.

When you downsize, you change the ratio between circumferential scale and jet width, the influence of boundary layers, and the relative size of mechanical imperfections.

A practical scaling strategy has three goals: preserve the dynamical regime (Rossby and Ekman numbers), preserve jet sharpness relative to circumference, and keep the system mechanically stable.

To preserve the dynamical regime, you want to keep Rossby number and Ekman number in the same ballpark when you change size. That means that if you reduce the jet radius, you may need to adjust differential forcing and background rotation correspondingly. Smaller systems often require higher rotation rates to maintain strong rotational control, because the timescales become shorter and boundary effects become proportionally larger. But higher rotation rates demand smoother mechanics, which is harder in a small DIY platform. This is the hidden trap: you can preserve the fluid dynamics only if you can preserve smooth rotation.

To preserve jet sharpness, you must avoid letting the jet become “too fat” relative to the ring circumference. In small tanks, a ring that is mechanically convenient may end up being a large fraction of the system size. That broadens the shear zone and destroys clean mode selection. The scaling principle here is that the ring width and the effective jet width should remain a small fraction of the jet radius. If downsizing forces you to violate this, you lose the polygon.

To preserve mechanical stability, you must recognize that wobble, tilt, and motor speed ripple do not scale down kindly. A millimeter of eccentricity is negligible in a large system but significant in a small one. A small imbalance creates proportionally larger accelerations. Thus, as you downsize, the tolerance on alignment becomes tighter. If you cannot improve alignment, downsizing makes success less likely even if the target Rossby and Ekman numbers look plausible.

A workable downsizing strategy is therefore not “shrink everything.” It is “shrink geometry modestly, but increase rotation control.” If you must use a small tank, you compensate by building a better lid, using a smoother motor controller, and paying more attention to leveling. You also accept that the hexagon may be harder to see, so you invest in visualization: strong lighting, high-contrast dye, and top-down video recording.

The most important conceptual point is that scaling is not primarily about matching dimensions. It is about matching regimes. A Saturn-like hexagon is a regime phenomenon: a rotating annular jet that selects a low-mode wave and locks into it. If you downsize and preserve the regime, the polygon survives. If you downsize and drift into a different regime—too viscous, too noisy, too broad a jet, too weak a rotational constraint—the pattern disappears or is replaced by a different kind of polygon that is not jet-based.

In other words, the hexagon is not fragile, but it is specific. Scaling laws exist to help you keep specificity while changing size. The builder's job is to use those laws the way engineers use them: as a guide to what must remain invariant, so that the thing you care about stays the same even when the apparatus changes.

5. Bill of Materials and Tools

A DIY annular-jet hexagon is not a "parts-heavy" project, but it is unusually sensitive to two kinds of quality: rotational smoothness and concentric alignment. That means your bill of materials should be biased toward components that reduce wobble, suppress slosh, and maintain stable speeds. The best way to control cost is not to buy the cheapest version of everything; it is to spend on the few elements that actually determine whether the jet becomes a clean, phase-locked polygon.

This section is written as a menu rather than a single prescription. There are multiple ways to build each subsystem. The goal is to give you options from budget to robust, while keeping the experimental intent intact: a steady rotating frame plus a separate, controlled differential ring drive that creates a narrow annular jet.

5.1 Tank and rotating platform options (budget to robust)

The tank must be round, stiff enough not to deform, and transparent if you want easy top-down observation. The platform must rotate smoothly under load and allow you to dial Ω repeatably.

Budget tank options typically include clear plastic storage containers, acrylic cylinders, or large clear bowls with near-circular symmetry. The main hazards here are subtle: many consumer containers are not truly round, have uneven wall thickness, or flex under load, which introduces asymmetry that the jet will "feel." If you go this route, prioritize a container that is visibly cylindrical, with uniform walls, and that sits flat without rocking.

A mid-tier tank option is a clear acrylic cylinder or polycarbonate tube segment mounted to a rigid base. This can be surprisingly affordable if you buy a short length of tube and fabricate a base plate. The advantage is consistent geometry. The disadvantage is that you need basic fabrication tools to make a watertight base and ensure the cylinder is square to the base.

A robust tank option is a purpose-built cylindrical vessel or a small clear acrylic aquarium cylinder. These are more expensive but reduce the time you spend fighting unintended asymmetries. If you are serious about repeatability and want to document mode transitions cleanly, the robustness is worth it.

For rotating platforms, the budget tier includes repurposed record turntables or lazy Susan assemblies driven by a motor. These can work, but they often lack torque and speed stability, especially under the mass of water. They also often have speed ripple that becomes an unwanted periodic forcing.

The mid-tier platform is a pottery wheel. Pottery wheels are built to carry heavy loads smoothly, provide variable speed, and tolerate some imbalance. They are one of the most “DIY-friendly” ways to get stable Ω without engineering a full rotation system from scratch. The primary downside is size and the need to manage splash and guarding.

The robust tier is a lab-grade rotating table or a heavy-duty turntable designed for filming or industrial use. This is overkill for many builds, but if you already have access to such equipment, it gives you the smoothest possible rotation and reduces the number of confounders.

One practical note that applies to all tiers: the diameter of the platform should exceed the tank base, and the platform should be rigid. Flexure introduces wobble. Wobble introduces spurious forcing. In this experiment, spurious forcing is not just an annoyance; it can be the difference between a locked hexagon and a wandering, incoherent wave.

5.2 Ring assembly options: bearing, seal, and drive coupling

The ring assembly is the signature feature of the Saturn-like build. It is where most DIY creativity goes, and it is also where most failures originate. You want a ring that rotates smoothly relative to the co-rotating lid, remains concentric with the tank axis, does not chatter, and couples to the water without turning into a turbulence generator.

The core of the ring assembly is a bearing interface. The simplest approach is to mount the ring on a low-friction bearing so it can rotate relative to the lid. A large-diameter lazy Susan bearing can be adapted for this role, especially if mounted above the lid and connected to the ring by a rigid annular support. The advantage is availability. The downside is that many lazy Susan bearings have noticeable play. That play becomes radial wobble at the ring, which broadens the jet and destabilizes the pattern.

A more precise approach is to use a ring supported by multiple small bearings arranged in a circle, so the ring rides on them like a rail. This reduces wobble if the geometry is built carefully. It also lets you keep the ring closer to the water surface without bulky bearing hardware. The tradeoff is more fabrication and alignment effort.

A robust approach is to use a proper rotary bearing assembly or a high-quality turntable bearing with minimal runout. If you can obtain a bearing with low axial and radial play, you get a direct improvement in the clarity and persistence of the polygonal mode.

Sealing is not always required, depending on whether your ring penetrates the lid and whether water can splash into the bearing region. But splash management matters. Water droplets in a bearing change friction unpredictably and can add noise. The easiest “seal” is simply a splash shield and careful lid geometry that prevents water from reaching moving parts. If the ring dips into the water, you may want a flexible skirt or a simple wipe seal to reduce splashing through gaps. The point is not watertight perfection; the point is stable friction and stable geometry run to run.

Drive coupling is how you rotate the ring at ω relative to the rotating frame. The simplest coupling is friction: a rubber wheel driven by a small motor presses against the ring’s outer edge. This is attractive because it is easy to build and tolerates small misalignment. It also acts as a slip clutch, which can prevent sudden shocks. The downside is that friction drives can introduce speed ripple and micro-slip, which may cause the polygon to jitter.

A belt drive is smoother and can reduce ripple. A small belt around the ring circumference, driven by a pulley, yields steadier ω if tension is consistent. The downside is that belts add compliance; the system can oscillate if not damped.

A gear drive is precise but usually noisy and can inject high-frequency vibrations. In this experiment, high-frequency vibration is rarely your friend. If you choose gears, you should treat vibration control as part of the design rather than as an afterthought.

A magnetic coupling is elegant: you drive a ring of magnets above the lid that couples to a magnetized ring or ferromagnetic strip below, transmitting torque through the lid without a shaft. This can produce very smooth forcing and keeps mechanics dry. It is more complex to implement but can pay off in stability.

In all cases, the practical goal is the same: stable ω with minimal wobble and minimal vibration.

5.3 Motors, controllers, power, and wiring on a rotating frame

You have two independent motions to control: Ω for the platform and ω for the differential ring. Many builders mistakenly try to drive both from one motor. That tends to collapse the experiment into a single swirl system. You want two controls, even if both are crude, because the entire tuning strategy depends on adjusting differential forcing without altering the global rotating frame.

For the platform, if you are using a pottery wheel or a turntable with built-in speed control, you may not need to modify the motor system at all. Your job becomes measurement and repeatability rather than motor design. If you are building your own platform, you want a motor with sufficient torque and a controller that provides smooth speed control under varying load. A gear motor can help by trading speed for torque and reducing sensitivity to load changes.

For the ring, a small DC gear motor is usually ideal. The ring does not require much torque if friction is low, but it requires smooth, controllable speed. Gear motors often provide stable low rpm operation without stalling. You want the ability to make fine adjustments, because the difference between m equals 7 and m equals 6 can be a small change in ω.

Controllers matter more than many DIY builders expect. A cheap pulse-width modulation controller can work, but some introduce audible or mechanical ripple. A controller that allows fine control at low speed and does not produce large torque pulses is worth seeking. If your ring drive visibly accelerates and decelerates on each cycle, you are effectively applying a periodic forcing that competes with the natural wave modes.

Power and wiring on a rotating frame is the second place DIY builds often break. If your ring motor is mounted on the rotating platform, it rotates with the tank. That is good because it avoids relative motion in the wiring for that subsystem. You can power it with a small battery mounted on the rotating frame, which eliminates slip rings and reduces noise. A battery plus a controller plus the motor can be a clean “self-contained” ring drive.

If you want longer runs or more consistent voltage, you might consider a slip ring. Slip rings add complexity and can introduce electrical noise, but they allow continuous power without battery management. For a first build, battery power is simpler and safer.

Regardless of approach, treat water and electricity with respect. Use low voltage. Keep electronics elevated and shielded. Provide drip loops in cables. Use splash guards. The experiment becomes much more enjoyable when you are not worried about a spill ending your controller.

A final note: mounting. Motors must be mounted rigidly. A motor that shifts slightly under torque is a vibration source. Vibration is an unwanted wave generator. In this experiment, mechanical rigidity is part of your fluid dynamics.

5.4 Low-cost measurement tools: tachometer, strobe, smartphone video

You can do this experiment by eye, but you will get much farther if you can measure Ω and ω even approximately. The patterns are sensitive, and repeatability requires being able to return to known conditions. Luckily, measurement can be cheap.

A simple optical tachometer is one of the highest value tools you can use. With a strip of reflective tape on the rotating platform and on the ring, you can measure rpm quickly and non-invasively. This lets you record the conditions that produced m equals 6 and return to them later. Even if absolute accuracy is modest, consistency is what matters.

A strobe light can be used as a qualitative measurement and diagnostic. By flashing at a known frequency, you can make the platform or the polygon appear stationary in your video or to your

eye. This is especially useful for determining whether the polygon is phase-locked and for measuring the pattern's drift rate relative to the rotating frame. A strobe does not have to be fancy; even some smartphone apps can produce usable strobe effects, though dedicated strobes are easier to tune precisely.

Smartphone video is the most powerful measurement device you already own. Top-down video lets you slow-motion the transition from circular jet to polygonal wave. It lets you count lobes reliably. It lets you estimate the pattern's rotation rate by tracking a corner frame-to-frame. It also reveals things you might miss in real time, such as subtle wobble, micro-slips in the ring drive, or slow drift in the jet radius.

A helpful practice is to place a simple angular reference in the field of view, such as a mark on the lid or a fixed pointer attached to the rotating frame, and then record both the reference and the polygon. This makes it easier to see whether the polygon is locked in the rotating frame or drifting relative to it.

Other low-cost tools include a bubble level for leveling the platform, a ruler or calipers for measuring ring radius and width, and a kitchen scale if you want to estimate the mass and balance of the rotating assembly. None of these are glamorous, but each reduces a confounder.

The overarching principle is that measurement is not about precision for its own sake. It is about making the experiment learnable. A DIY annular jet hexagon is a system with discrete attractors. The fastest way to reach the hexagon is to treat it like a tunable instrument, record what works, and return to those settings with confidence. Measurement turns your build from "I saw it once" into "I can reproduce it," and reproducibility is the difference between a demo and a real laboratory analogue.

6. Mechanical Build: A Repeatable Annular Ring Drive

If the Saturn-like hexagon is a regime phenomenon, then the mechanical build is what determines whether you can access that regime reliably. The physics may want to give you a clean polygonal jet, but the apparatus can easily sabotage it by adding unwanted forcing: wobble, vibration, surface slosh, and slowly drifting misalignment. A successful build is therefore not judged by elegance or cost. It is judged by whether it produces a stable, repeatable annular jet that can be tuned through mode numbers without constantly collapsing into noise.

The annular ring drive is the heart of that repeatability. Your goal is a ring that can rotate at a controlled differential rate relative to the tank, while the entire assembly rotates together as a single rigid body. This is a subtle demand. It means the ring drive motor, controller, battery (if

used), and mechanical coupling all rotate with the tank. From the ring's perspective, it is rotating relative to the lid; from the outside world's perspective, everything is spinning. The simplest way to keep this from becoming a mess is to build in layers: a rigid co-rotating lid as the structural reference, a low-friction ring mount that preserves concentricity, and a drive mechanism that transmits torque smoothly without introducing ripple.

The following subsections describe not only how to build the parts, but what each part is protecting you from. In this experiment, every mechanical choice is also a dynamical choice.

6.1 Building the co-rotating lid to suppress surface slosh

The lid does two jobs simultaneously: it stabilizes the fluid surface and it provides a structural platform for the ring drive. Both are essential.

From a fluid standpoint, surface slosh is your enemy because it introduces gravity waves and irregular oscillations that compete with the jet's wave selection process. The annular jet hexagon is a low-frequency, large-scale pattern that emerges from a delicate instability and saturates into a locked state. Surface waves inject additional modes, stir the fluid vertically, and blur the jet. A lid reduces the free surface's ability to support those waves.

From a mechanical standpoint, the lid is a rigid reference plane. The ring must remain concentric and stable. If you mount the ring on the tank rim alone, the rim can flex and the ring can wobble. A well-built lid ties the geometry together and reduces deflection.

A practical co-rotating lid can be built from a flat, rigid sheet—acrylic, polycarbonate, or plywood sealed for splash resistance. Acrylic and polycarbonate are attractive because you can see through them, which makes top-down observation easier. Thickness matters. A thin sheet can flex and introduce oscillations. A thicker sheet is stiffer and acts more like a true reference plane. In a DIY context, stiffness is worth more than transparency, but if you can have both, it is ideal.

The lid should fit the tank closely enough to rotate as a unit with it. There are multiple ways to achieve this. The simplest is to mechanically attach the lid to the tank or to a rigid frame that holds the tank. Another approach is to create a lid that rests on the tank rim and is clamped or taped so it cannot slip. The key is that the lid must not lag behind the tank or twist relative to it; any relative motion introduces a time-varying boundary condition at the surface.

You also want a small gap between lid and water surface. If the lid is too high, it does not suppress slosh effectively. If it is too low, it may contact the water during transients or at higher speeds, causing unpredictable drag and splashing. In practice, you aim for a consistent clearance that is small enough to damp waves but large enough to avoid contact. A simple way to control clearance is to place small spacers at the rim so the lid height is repeatable from run to run.

You may need one or more access ports in the lid: a small hole for dye injection, a hole for a pipette or tube, and possibly a larger removable plug for cleaning. These holes should be as small as practical to avoid reintroducing wave motion through airflow and to preserve structural stiffness. If you cut holes, reinforce around them, because holes concentrate stress and can encourage flexing.

A well-built lid produces an immediate behavioral change in the system: during spin-up, the water settles faster, the surface looks calmer, and dye injected near radius R forms a sharper ring rather than rapidly dispersing into a mottled pattern.

6.2 Mounting a concentric ring with low wobble and low friction

The ring mount is where you either create a clean annular jet or create persistent frustration. The ring must satisfy three requirements at once: it must be concentric with the tank axis, it must rotate smoothly with low friction, and it must maintain its alignment over time without developing wobble as speed changes.

Concentricity is non-negotiable because the jet you create is an annulus centered on the rotation axis. If the ring's center is offset from the tank's center even slightly, the shear forcing becomes stronger in one sector than another. This introduces an artificial "preferred corner." The system may still form a polygon, but it will be biased and it may drift unpredictably. Worse, it may lock into patterns that are more reflective of mechanical eccentricity than of fluid dynamics.

Low friction is important because friction variations translate into torque variations, which translate into speed ripple. Speed ripple is effectively periodic forcing. Periodic forcing can entrain the wave pattern at an external frequency or prevent it from locking at all.

Low wobble is important because wobble broadens the effective jet radius. If the ring radius oscillates, the shear layer is not anchored at one radius. The jet becomes smeared, the mode competition increases, and the polygon becomes unstable.

A common DIY solution is to use a large bearing between the ring and the lid. A lazy Susan bearing can work if it is of decent quality and if you reduce play. But many such bearings have noticeable clearance. That clearance becomes wobble. If you use this approach, you can sometimes improve it by preloading the bearing slightly so it does not rattle. Preloading means applying gentle axial force so the bearing surfaces remain in contact consistently. However, preload also increases friction. The tradeoff must be managed.

A more controlled DIY approach is the "roller bearing ring": mount several small bearings (for example, skate bearings) around a circular track so the ring rides on them. This can yield very low friction and low wobble if the geometry is built carefully. The ring becomes a rigid hoop that

rolls on the bearings. You can add a second set of bearings on the side to constrain lateral motion. The advantage is that you can tune the preload and constrain wobble strongly. The disadvantage is that it requires precise placement of multiple bearings.

In either approach, the ring itself should be rigid. A flexible ring is a vibration source and can deform under drive torque, especially with a friction wheel drive. A rigid ring, even if it is slightly heavier, maintains its geometry. You can build the ring from a stiff plastic hoop, an acrylic annulus, or a lightweight metal ring. The important property is that it does not warp and that it can be mounted securely to its bearing interface.

Coupling to the fluid should be gentle. The ring can dip slightly into the water or present a narrow edge that drags the surface layer. If it dips too deeply, it becomes a paddle and injects turbulence. If it barely touches, it may not create a strong enough jet. A practical compromise is a small immersion depth that creates shear primarily in a thin surface layer, allowing the rotating background to distribute the effect downward without violent stirring.

Your “success test” for the ring mount is simple: when the tank is not spinning, you should be able to rotate the ring smoothly by hand with minimal resistance and no visible wobble. When the tank is spinning, the ring should maintain its radius and not visibly oscillate.

6.3 Drive strategies: belt, friction wheel, gear, or magnetic coupling

The drive strategy is how you impose ω cleanly. The choice here is not primarily about efficiency; it is about smooth torque delivery and minimal vibration. You want the ring to rotate at a stable differential rate, and you want that rate to be adjustable in small increments.

A friction wheel drive is the simplest. You mount a small rubber wheel on a motor shaft and press it gently against the ring’s edge. The wheel transmits torque through friction. The best feature of this approach is that it naturally slips if the ring jams, reducing shock loads. It also lets you adjust ω by adjusting motor speed easily. The main drawback is micro-slip and torque ripple, especially if the ring edge is not perfectly uniform. Micro-slip appears as small fluctuations in ring speed. Those fluctuations can show up as jitter in the polygon’s rotation rate. If you use friction drive, you improve it by ensuring the ring edge is smooth, the wheel has consistent contact pressure, and the motor controller is stable at low speeds.

A belt drive is smoother in many cases. You put a belt around the ring or around a pulley attached to the ring, then drive it with a motor pulley. Belts reduce the high-frequency ripple that friction drives can introduce, and they can maintain stable speed if tension is consistent. The downside is complexity: belt tension must be managed, and belts introduce compliance. If the system is lightly damped, that compliance can lead to oscillations in speed. In practice, belts tend to work best when the ring rotates slowly and the belt is properly tensioned so it does not slip.

A gear drive provides precise ratio control, but it often introduces vibration. Gears can chatter, especially if alignment is imperfect. The ring and gear interface must be manufactured with decent tolerance, or you get periodic forcing at gear mesh frequency. That forcing can compete with the natural wave pattern. Gears also tend to be louder and more mechanically "sharp" than belts or friction wheels. In this experiment, sharpness is often the opposite of what you want. If you must use gears, choose large, low-tooth-count gearing that runs slowly and smoothly, and mount the motor rigidly with proper alignment.

A magnetic coupling is the most elegant if you can implement it. In this approach, the motor drives a ring of magnets above the lid. The ring itself has magnets or a ferromagnetic strip, and the torque transmits through the lid without physical contact. The advantage is smoothness and isolation: no shaft penetration, no water exposure, and often very low mechanical noise. The drawback is design effort. You must arrange magnets so that coupling is strong enough at the desired ω but not so strong that it creates "cogging" that produces its own periodic forcing. A well-designed magnetic coupling can yield very clean differential rotation and is particularly attractive if you care about long-run stability and minimal mechanical artifacts.

Regardless of strategy, the drive motor should be mounted on the rotating frame, not in the lab frame, to avoid tangled wires and to keep the forcing purely differential in the rotating coordinate system. The motor should also be mounted with vibration isolation if necessary, but isolation should not allow the motor to shift under torque. The best vibration isolation is often rigid mounting plus balanced components rather than soft mounts that allow wobble.

A practical tuning note: you want the ability to change ω in small steps. This argues for a controller that has fine control in the low-rpm regime. The difference between stable m equals 7 and stable m equals 6 can be subtle, and coarse control makes it hard to settle into the right attractor.

6.4 Alignment and leveling: how small errors kill symmetry

In an experiment built on rotational symmetry, small mechanical errors are not small. They do not simply add noise; they introduce preferred directions. Preferred directions break the very condition that allows the system to select a clean integer mode on its own.

Alignment errors come in three main forms: tilt, eccentricity, and runout.

Tilt is when the rotation axis is not vertical. Even a small tilt introduces a gravity component that varies around the circle in the rotating frame. This produces a periodic forcing that can mimic or disrupt the natural wave modes. Tilt can also cause the fluid depth to vary slightly with angle, which changes local dynamics and can bias where lobes form.

Eccentricity is when the tank or ring is not centered on the rotation axis. This produces a repeating pattern of stronger and weaker forcing as the system rotates. Eccentricity is especially damaging for annular jets because it smears the jet and can force the system into patterns that are locked to the mechanical offset rather than to the fluid's natural mode selection.

Runout is when the ring or platform wobbles up and down as it rotates. This introduces vertical oscillations that can couple into the fluid as pressure variations and surface waves, precisely what the lid is supposed to suppress. Runout also makes the ring immersion depth vary, causing periodic changes in forcing strength.

The reason these errors "kill symmetry" is that they create an external signal at exactly the frequency the system is sensitive to: the rotation frequency and its harmonics. The polygonal pattern you are trying to produce is itself a low-mode structure that lives on the ring. If your apparatus injects a strong periodic component, it will compete with or dominate the fluid's intrinsic selection. You may still see a polygon, but it will not be a clean demonstration of the dynamical phenomenon. It will be partly a driven response.

This is why leveling and centering are worth real effort. Use a bubble level. Shim the platform. Mark the center and ensure the tank sits exactly there. Add balancing mass if the tank assembly is heavy on one side. Check ring concentricity by measuring the gap between ring and tank wall at multiple angles. If the gap changes, you have eccentricity. Fix it before you chase fluid parameters.

A simple diagnostic is to spin the empty assembly and observe from the side and from above. If you see wobble in the rim, you will see it amplified in the fluid. Another diagnostic is to place a pointer near the ring edge and watch the gap as it rotates. A stable gap indicates low runout; a varying gap indicates eccentricity or wobble.

In a DIY context, it is tempting to accept small errors and "let the fluid sort it out." But the entire appeal of the Saturn-like hexagon is that it is a self-organized pattern. If you allow mechanical asymmetry to do the organizing, you lose the point. The build's purpose is therefore to remove external biases until the only remaining bias is the physics itself: the interaction of rotation, shear, and instability.

When you treat alignment and leveling as first-class parts of the experiment, the rest becomes dramatically easier. The jet becomes narrower, the onset of polygonal modes becomes cleaner, and the pattern becomes more stable. The build stops feeling like a contraption and starts feeling like an instrument. That is the moment where a DIY apparatus becomes a laboratory analogue rather than a rotating art project.

7. Electrical and Safety Architecture

A rotating-tank experiment is a rare combination of hazards: moving machinery, a large mass of water, and wiring that may be forced to move with the system. The good news is that you can design the electrical and safety architecture so that these hazards are mostly neutralized. The bad news is that if you ignore this layer, you will either have an unsafe build or an experiment so fragile you will avoid running it.

In a Saturn-like annular jet build, the electrical system has two jobs: provide stable speed control for ω (and sometimes Ω if your platform is DIY), and do so in a way that remains safe under splashes, vibration, and rotation. The safety system has one job: keep rotating mass and electricity from interacting with bodies, hair, loose clothing, and water in uncontrolled ways.

If you get this layer right, the experiment becomes something you can run repeatedly, confidently, and alone. If you get it wrong, the experiment becomes something you run once while holding your breath and hoping nothing goes sideways.

7.1 Why low-voltage DC is the right choice

Low-voltage DC is the default choice for this experiment for three reasons: shock risk, controllability, and mechanical simplicity.

Shock risk is straightforward. Water and mains voltage do not belong in the same DIY apparatus unless you are building to professional standards with proper enclosures, grounding, and residual-current protection. Even if you are confident in your wiring, splashes and condensation are chaotic. Low-voltage DC drastically reduces the consequences of a mistake. It is not “risk free,” but it moves you out of the most dangerous regime.

Controllability is the second reason. The differential ring drive is a low-speed application. You want ω in the realm of fractions of a rotation per second, and you want fine adjustments. DC gear motors are excellent at producing stable low rpm with reasonable torque. DC motor controllers also tend to be compact and easy to integrate on the rotating frame. You can mount the entire system—motor, controller, battery—on the rotating platform and treat it as a self-contained module.

Mechanical simplicity is the third reason. The moment you put an AC-powered device on a rotating platform, you must solve for cable management, strain relief, and slip or twist tolerance. A DC system can be battery powered, eliminating the need for moving power cables entirely. The mechanical build becomes cleaner and the failure modes shrink dramatically.

There is also a psychological reason worth acknowledging: if the system feels safe, you will run it more. You will tune it. You will learn it. A safe design increases the experiment’s “iteration rate,” and iteration rate is what gets you from “interesting” to “reproducible.”

7.2 Power on a rotating platform: battery versus slip ring

The key architectural decision is whether you bring power onto the rotating platform continuously or carry power on the platform itself.

Battery power is usually the best choice for early builds. The logic is simple: if the ring motor and controller rotate with the platform, then a battery rotating with them provides power with no moving electrical interface. No moving interface means fewer failure modes. A battery also isolates the ring drive from external electrical noise, which can show up as speed ripple. In practice, a battery pack sized for modest current draw can run a session long enough to tune a stable polygon.

Battery power has two costs. The first is voltage drift as the battery discharges. The second is mass and balance. Voltage drift can be mitigated by using a controller that regulates motor speed reasonably well, or by running sessions short enough that drift is negligible. Mass and balance can be handled by mounting the battery close to the axis and counterbalancing if necessary.

Slip rings are the professional solution for continuous power transfer to rotating systems. A slip ring provides electrical contacts through a rotating interface, allowing the platform to spin indefinitely while receiving power from the stationary frame. Slip rings are useful if you want long runs, higher power, or if you want to power multiple subsystems without swapping batteries.

Slip rings also introduce complexity. They can introduce electrical noise and intermittent contact, both of which can create speed ripple in the motor. They add points of failure: contacts can wear, oxidation can increase resistance, and mechanical alignment must be maintained. They also require careful strain relief, because a rotating platform with a slip ring still has wires that must be managed.

A practical rule is to start with battery power unless you know you need a slip ring. Battery power gets you into the regime faster. Once you have demonstrated stable polygons and you want longer sessions, continuous logging, or more complex instrumentation, you can consider a slip ring upgrade.

One hybrid approach is worth mentioning: keep the ring motor on battery, but keep any optional sensors or lights on separate power. This isolates the drive from noise and keeps your main success variable clean.

7.3 Splash protection, guarding, and safe operating practices

Safety in this build is not a single feature. It is a system. The three pillars are splash protection, guarding, and operational discipline.

Splash protection is essential because water droplets will find the worst possible place to land. You should assume splashes will happen during spin-up, during mode transitions, or when you overdrive ω. You should also assume that fine mist can occur, especially at higher Ω, even if you do not see dramatic splashing.

The most effective splash protection is enclosure by geometry. A lid reduces surface motion and therefore reduces splashes at the source. A raised rim or skirt around the lid can catch droplets. A clear shield around the platform perimeter can prevent outward-thrown droplets from reaching the motor or electronics. You do not need to build a sealed box; you need to design surfaces that direct water away from electrical components.

Where electronics must exist on the rotating platform, they should be elevated, mounted in simple enclosures, and placed as close to the axis as possible. Closeness to the axis reduces centrifugal exposure and reduces the chance that a droplet thrown from the rim reaches the component. Even a simple plastic box with a lid, with cable pass-throughs oriented downward, significantly improves resilience. The goal is not waterproofing in the submarine sense; it is drip control.

Guarding addresses rotating mass hazards. A tank full of water on a spinning platform has real momentum. If something catches—hair, sleeves, a lanyard—the outcome is unpredictable. The safest practice is to create a physical barrier between you and the rotating system during operation. A low guard ring around the platform can prevent accidental contact. A larger transparent shield can protect against both contact and splashes.

Operational discipline is the human layer. You do not lean over a spinning system. You do not reach across it. You do not adjust the ring drive while your hand is near the ring edge. You stop the platform before making mechanical changes. You keep loose clothing, jewelry, and long hair secured. You treat the system the way you would treat a lathe or any rotating machinery: with respect for the fact that it does not care about your intentions.

There is also a safety practice specific to the differential ring: the ring is often near the lid edge, which is where hands naturally go when adjusting things. That is exactly where you should not put fingers during operation. If you need to adjust contact pressure on a friction wheel, build an adjustment knob or lever that can be operated away from the ring edge, or stop the system first. Small design choices like remote-adjustment screws are safety features, not luxuries.

7.4 A pre-run checklist to prevent the common failures

A pre-run checklist is the fastest way to prevent both safety incidents and wasted runs. It turns your build into an instrument with a consistent start state. The checklist below is deliberately short but covers the failure modes that appear most often in DIY annular jet builds.

7.4.1 Mechanical integrity checks

Confirm the tank is firmly attached or seated and cannot drift outward on the platform.
Confirm the platform rotation is smooth with no visible wobble.
Confirm the lid is secure and co-rotates with the tank without slipping.
Confirm ring concentricity by checking the ring-to-wall gap at multiple angles.
Spin the ring by hand and confirm it rotates smoothly without binding.
Confirm the drive mechanism contacts the ring consistently without chatter.

7.4.2 Leveling and symmetry checks

Confirm the platform is level using a bubble level in two perpendicular directions.
Confirm the tank is centered on the platform axis.
Confirm any heavy components on the rotating frame are balanced or placed near the axis.

7.4.3 Electrical checks

Confirm all electrical components are enclosed or shielded from splashes.
Confirm cable routing is secured with strain relief and cannot contact rotating edges.
Confirm the battery is securely mounted and cannot shift under rotation.
Confirm the controller's adjustment knobs are accessible without reaching over the tank.
Test the ring motor briefly at low speed and verify stable operation without stutter.

7.4.4 Splash and guard checks

Confirm a splash guard or perimeter shield is in place if you plan to increase Ω.
Confirm the lid clearance is correct: close enough to damp slosh, not so close it contacts the water.
Confirm the work area is clear of items that could be pulled into the rotating system.

7.4.5 Operational readiness checks

Tie back hair and remove loose jewelry or dangling clothing.
Confirm you have a clear stop procedure: how to cut power quickly to the platform and ring drive.
Confirm you have dye injection tools ready so you do not improvise mid-run.
Start with conservative Ω and ω settings; never jump directly to high speed.

This checklist prevents the two most discouraging outcomes: "I got splashes on my electronics" and "I wasted an hour chasing a pattern that was actually caused by wobble." It also prevents the most serious outcome: a preventable injury.

The deeper point is that safety and experimental quality are aligned. The same features that make the build safer—rigid mounting, stable wiring, shields, controlled access—also make the fluid dynamics cleaner. A safe build is usually a better experiment. In a project whose goal is a

delicate, phase-locked pattern, designing for safety is not a separate concern. It is part of the physics.

8. Fluid Preparation and Visualization

The experiment lives or dies by what you can see. An annular jet hexagon is not just a shape; it is a shear layer that becomes a wave and then locks into a stable mode. If your visualization is poor, you will misdiagnose what is happening and tune in the wrong direction. If your visualization is good, you can watch the jet form, watch the instability onset, and watch the mode settle. That feedback loop is what makes a DIY build learnable.

Fluid preparation is also an experimental control knob in disguise. The fluid's viscosity, the way dye spreads, and the way tracers move all affect whether you get a crisp jet, a smeared shear zone, or turbulence. In a rotating system, you want enough damping to suppress irrelevant small-scale mess, but not so much that the jet is erased. Visualization choices can amplify or hide those differences.

This section describes the practical choices that make the jet visible and interpretable, from simple food coloring to fluorescent dyes and low-effort quantitative tracking.

8.1 Water versus water–glycerin mixes: choosing viscosity

Water is the default fluid for good reasons: it is cheap, clean, and predictable. It has low viscosity, which means the flow can remain lively and develop coherent waves when rotation is strong enough. It also means that boundary layers are relatively thin and that the bulk can behave in a quasi-two-dimensional way when Ω is high.

But water has a common DIY failure mode: it can be too lively. If your platform and ring drive are not exceptionally smooth, small mechanical disturbances can trigger bursts of turbulence that smear the jet and break phase-locking. When that happens, adding a controlled amount of viscosity is a legitimate fix, not a hack. Water–glycerin mixtures are the easiest way to adjust viscosity without changing other properties too dramatically.

The purpose of adding glycerin is not to "thicken the fluid so it makes shapes." The purpose is to shift the balance between coherent wave growth and small-scale noise. A modest increase in viscosity can reduce chatter, slow down the development of unwanted small eddies, and make the annular jet sharper and more stable. It can also make spin-up slower, so you need more patience before the flow reaches steady behavior.

A practical approach is to treat viscosity as a secondary tuning knob. Start with pure water. If you can form a clean jet but the polygon never locks, or if the flow repeatedly breaks into small-

scale turbulence when you cross the instability threshold, then consider adding glycerin. If your jet never forms cleanly and dye immediately diffuses into a broad smear, then increasing viscosity may make the jet even harder to see; in that case you likely need more rotation discipline, better slosh suppression, or a narrower forcing ring rather than a thicker fluid.

If you do use glycerin, consistency matters. Mix thoroughly before runs. Temperature matters too: viscosity changes with temperature, and so does the behavior of boundary layers. If you are chasing reproducibility, keep the room temperature roughly stable and note whether the fluid feels warmer after extended runs.

A useful mindset is this: water is the “high-performance” fluid that rewards good mechanics; water–glycerin is the “forgiving” fluid that can make a marginal mechanical build workable. Neither is universally superior. Choose based on what your system needs.

8.2 Dye injection methods that reveal the annular jet cleanly

A Saturn-like annular jet experiment is easiest to interpret if the jet is visible as a thin ring before it becomes polygonal. The visualization goal is therefore a narrow dye band centered at radius R that tracks the jet. That band should persist long enough for you to watch it develop undulations and then lock into a stable mode.

The most common mistake is to pour dye into the tank. That almost guarantees that dye will mix broadly, obscuring the jet. The right approach is controlled injection.

The simplest controlled method is syringe injection through a small port in the lid. You use a syringe with a length of thin tubing so you can place the injection point precisely at radius R near the surface. You inject a small volume slowly, aiming to seed the shear layer rather than to tint the whole tank. You can also inject multiple small pulses at different angles, which helps reveal whether the jet is truly annular and whether it has a consistent radius.

A second method is “line injection”: use a thin tube or a small nozzle that slowly releases dye while you rotate the tank at steady Ω and steady ω. This creates a continuous dyed ring. The advantage is clarity: you get a full annulus quickly. The disadvantage is that you are adding dye for longer and may start to contaminate the surrounding region if you overdo it.

A third method is to use a weak dye and inject upstream of where you want to observe. In a rotating frame, the meaning of “upstream” is subtle, but the principle is that you do not need to inject exactly on the jet centerline. Injecting slightly inside or outside can reveal the shear direction and help you identify where the jet core is relative to the ring.

A practical tip is to use very small amounts. Many dyes are far more concentrated than needed. The goal is contrast, not saturation. If you can still see white water behind the dye band, you are in the right regime. If the whole tank becomes colored, you have lost your ability to interpret.

Another useful practice is to use two colors: inject one color slightly inside radius R and another slightly outside. This makes the shear layer visible as the interface between the two. It can be surprisingly informative when you are diagnosing whether the “jet” is actually narrow or whether you have created a broad region of differential rotation.

Injection depth matters as well. If you inject too deep, you are visualizing subsurface motion that may not correspond cleanly to the surface jet. If you inject at the surface, you see the jet more directly but may also be more sensitive to surface effects. In most builds, a near-surface injection is best because the ring forcing is also near the surface. You can always experiment with deeper injections once the basic behavior is reliable.

8.3 Tracers and lighting: maximizing contrast for top-down imaging

Dye shows stream structure, but tracers show motion. A good DIY setup uses both. Dye reveals the jet as a band; tracers reveal whether the flow is smooth, whether there are coherent vortices riding the jet, and whether turbulence is intruding.

For tracers, you want particles that are neutrally buoyant or nearly so, small enough to follow the flow, and visible against the background. Common household tracers like pepper or glitter can work, but they have drawbacks: pepper tends to clump and can be influenced by surface tension; glitter can be too heavy or reflect light unpredictably.

A better DIY tracer is very small, lightweight particles that remain suspended. Some people use fine hollow glass spheres or microbeads, but availability varies. If you are keeping the project simple, start with dye alone and add tracers only if you need motion cues to interpret what you are seeing.

Lighting is where you can gain a lot of visibility for little effort. The best top-down imaging is achieved when the surface is evenly lit and glare is controlled. Overhead lights often create reflections that hide subtle patterns. A ring light or diffuse side lighting can be superior because it reduces bright specular reflections off the water surface.

A simple improvement is to place a matte, high-contrast background under the tank. A dark background makes bright dye stand out; a light background makes dark dye stand out. Choose your dye color accordingly. If you use fluorescein later, the background choice becomes less important because fluorescence can overwhelm ambient reflections.

If you are recording with a smartphone, lock exposure if possible so the camera does not “hunt” as the bright dye moves. Also, mount the phone or camera rigidly above the tank. Handheld recording introduces motion that makes it hard to count lobes reliably.

Contrast is not only about color; it is also about sharpness. A calm surface yields a sharper image. This loops back to the lid. Slosh suppression is part of visualization. If the lid is doing its job, you will see the dye band as a crisp feature rather than a shimmering smear.

8.4 Optional upgrades: UV fluorescein, polarizers, simple PIV workflows

Once you can produce a stable jet and a polygonal mode, upgrades become valuable because they let you document the phenomenon in a more persuasive and even semi-quantitative way. None of these are required, but each can turn a DIY demo into something that looks and feels like a laboratory measurement.

UV fluorescein is the highest impact upgrade. Fluorescein dye, illuminated with blue or near-UV light, produces bright fluorescence that is easy to see even at low concentration. This allows you to inject extremely small amounts while maintaining high contrast. It also reduces the impact of background reflections because the signal is coming from fluorescence rather than from ambient light scattering. If you use fluorescein, you can often obtain a clear annular jet ring and polygonal undulations that are dramatically easier to record.

If you go this route, treat it as a visualization system: use a blue LED light source, use simple filtering if needed, and control ambient light. You do not need a professional fluorescence setup to get a strong benefit. The main caution is that fluorescein can stain surfaces, and you should protect your workspace accordingly.

Polarizers are a subtle upgrade. A polarizing filter on the camera lens can reduce glare from the water surface and improve visibility of dye patterns. This is especially useful if your lid is transparent and you have overhead lighting. Rotating the polarizer while viewing the surface can dramatically change the visibility of features. It is a cheap tool that often feels like magic the first time you use it in a reflective-water setting.

Simple PIV workflows are the bridge to quantitative flow analysis. Particle image velocimetry in its full form is complex, but you can adopt a “toy PIV” approach: seed the surface lightly with tracers, record top-down video, and use simple software to track particle motion between frames. Even without advanced algorithms, you can estimate the local velocity field in the jet region and measure the wave’s phase speed and amplitude.

In a DIY context, a minimal workflow is to record stable video with a known frame rate, then use a basic tracking tool to follow a few tracers around the jet. You can measure how fast the jet transports tracers relative to the rotating frame, estimate the polygon’s drift rate, and document how those change as you vary ω. This turns your narrative from “I saw a hexagon” into “the wave is phase-locked at a measurable rate and transitions between modes at repeatable settings.”

If you do pursue tracking, good lighting and high contrast matter even more. That is why fluorescein and glare control are not just aesthetic upgrades; they improve your signal-to-noise ratio for any measurement.

The overall philosophy of fluid preparation and visualization is this: first make the dynamics happen, then make them readable. The best visualization is the one that tells you what to change next. When you can see the jet as a narrow ring, see the instability as growing undulations, and see the final state as a locked polygon, you have an experimental loop. And once you have that loop, the hexagon becomes a matter of tuning rather than luck.

9. Operating Protocol: From Spin-Up to Hexagon

A Saturn-like annular jet hexagon is not something you “turn on” in one step. It is a sequence: establish a disciplined rotating fluid, introduce controlled shear at a chosen radius, let the annular jet sharpen, cross the instability threshold gently, and then tune until the system chooses the m = 6 attractor. The procedure matters because the fluid has memory. If you rush spin-up, you start with slosh. If you slam on differential rotation, you inject turbulence. If you chase polygons before the jet is established, you are tuning noise rather than dynamics.

This section gives a repeatable operating protocol that minimizes wasted runs and maximizes the probability that the first polygon you see is stable and interpretable. The goal is to make each run a clean experiment: you can adjust one variable at a time, observe the response, and converge toward a hexagon.

9.1 Spin-up to solid-body rotation: how long and how to tell

Spin-up is the process by which the fluid transitions from “lab frame water” to “rotating frame water.” In a successful run, you do not begin differential forcing until the bulk of the fluid is close to solid-body rotation with the tank. If you begin too early, you are effectively forcing a fluid that is still full of residual transients, and those transients will compete with the annular jet instability.

How long spin-up takes depends on your rotation rate, fluid depth, and viscosity. Instead of treating spin-up as a fixed time, treat it as a condition you verify.

A practical approach is:

Start the platform at a low Ω and ramp up smoothly to your target Ω over tens of seconds rather than instantly. Sudden acceleration produces large transient flows and surface motion. A smooth ramp is calmer and reduces the time needed for the system to settle.

Once at target Ω, wait. In many DIY builds, several minutes is a realistic spin-up time. If your fluid is deeper or more viscous, you may need longer. If you are in a rush, you will usually pay for it later by fighting incoherent patterns.

How to tell when you are close to solid-body rotation is more important than the exact time.

One clear indicator is surface calmness. With a good co-rotating lid, the surface should look quiet rather than oscillatory. If you still see rhythmic sloshing or obvious wave trains, you are not ready.

A second indicator is tracer behavior. If you sprinkle a few tracers near the surface, in a fully spun-up system they will rotate smoothly with minimal relative motion. If tracers show obvious spiraling or radial drift not tied to your intended jet radius, the system is still adjusting.

A third indicator is dye behavior. Inject a tiny amount of dye near a radius that is not your intended jet radius. In a spun-up system, that dye will tend to form a smooth arc and wrap around without chaotic mixing. If it immediately breaks into streaks, folds, and rapid dispersion, the bulk flow still has transient shear.

You can also use a simple “stop-and-see” test if you are diagnosing a stubborn build: briefly reduce Ω slightly and return. If the fluid responds dramatically, you were still in a transient state. If the response is mild and quickly damped, you are closer to steady rotation.

The basic discipline is: do not begin differential forcing until the fluid looks boring. Boring is good. Boring is the stable stage on which the annular jet instability can be seen cleanly.

9.2 Turning on differential rotation without triggering turbulence

The most common way to ruin a run is to turn on ω too aggressively. Differential rotation is the energy source for the instability, but it is also the easiest way to inject turbulence if applied as a shock rather than as a controlled perturbation.

The principle is to treat ω as a slow ramp, not a switch.

Begin with ω near zero. With the tank spun up, start the ring drive at the lowest setting where it rotates smoothly. Let it run there for a short period so the shear begins to develop gradually. Only then increase ω in small increments.

Small increments matter. In the m = 6 regime, you are often working with differential rates that are a small fraction of Ω. A change that looks “tiny” on your controller may represent a substantial change in forcing. Because the system has discrete mode attractors, you do not need large changes to move between them.

Between increments, wait long enough for the flow to respond. This is a frequent error: builders change ω, see no immediate result, change again, and then the system suddenly becomes turbulent. What happened is that the flow was responding on a longer timescale, and multiple changes stacked. A good practice is to change ω, wait, observe, and only then decide whether another change is needed.

Signs that you are triggering turbulence include:

The dye band broadens suddenly instead of remaining sharp.
Tracers begin to show small-scale chaotic motion rather than smooth advection.
The surface becomes visibly disturbed despite the lid.
The jet fails to maintain a clear radius and becomes a wide swirling region.

If you see these, the corrective action is usually not “push harder.” It is “back off and let it settle.” Reduce ω slightly, maintain Ω, and wait. Often the system can return to a cleaner state if you do not keep forcing.

A subtle but important point is that turbulence can be triggered mechanically as well as dynamically. If your ring drive has chatter, gear noise, or contact slip, the turbulence you see may be mechanically injected rather than a true fluid instability. That is why smooth motor control and stable mechanical contact are part of the operating protocol, not just the build.

9.3 Establishing the jet at radius R before chasing polygons

The annular jet must exist as a coherent structure before the polygon can form. This sounds obvious, but in practice it is easy to skip. Many runs fail because the builder starts counting corners before the jet is even a well-defined ring.

The first goal after turning on ω is to establish a narrow, stable jet at radius R. You should be able to see it as a sharp dyed annulus that remains narrow over multiple rotations. Only then do you pursue instability.

A practical sequence is:

Inject a small amount of dye at radius R, near the surface. Observe whether it forms a ring-like band that wraps around. If the dye immediately drifts inward or outward, your jet radius may not be where you think it is, or the flow has significant secondary circulation. Adjust ring immersion depth, lid clearance, or simply wait longer for the jet to develop.

If the dye band forms but is too broad, you likely have a broad shear zone rather than a narrow jet. This can happen if your ring is too wide, the ring couples too strongly (acting like a paddle), or ω is too high too soon. The remedy is often to reduce ω and ensure the ring is imposing shear

gently. Sometimes a modest increase in Ω helps tighten the jet by making the flow more geostrophic.

Once the dye band looks stable and narrow, look for signs of wave onset. The earliest sign is a gentle sinusoidal undulation of the dyed ring. At this stage, do not try to force it into m = 6. Your first objective is to confirm that the jet can support a coherent wave without breaking.

If you are not seeing any undulation even after the jet is established, you may still be below the instability threshold. Increase ω slowly until the ring begins to meander.

If you are seeing chaotic meanders rather than a coherent wave, the system may be too noisy. At this stage, it is often more effective to improve slosh suppression, reduce mechanical ripple, or slightly increase viscosity than to keep turning knobs aggressively.

The best runs are the ones where the progression is clean: circular dyed jet, gentle undulation, growing wave amplitude, then saturation into a stable polygon. If you skip the “circular dyed jet” stage, you are not chasing Saturn’s hexagon. You are chasing an uncontrolled swirl.

9.4 The tuning sequence that biases toward m = 6

Once you have a stable jet and a coherent wave onset, tuning becomes a structured process. The key is to bias the system toward the m = 6 attractor without overdriving it.

A successful tuning sequence typically looks like this:

First, fix Ω. Choose a background rotation rate that gives you a disciplined flow and a clear jet. Do not change Ω frequently during a run, because changing Ω changes both your rotational constraint and your effective frictional regime. In practice, you want Ω stable so the system has a stable “planet.”

Second, adjust ω slowly to enter the polygonal regime. As you increase ω, the wave amplitude grows. At some point the wave begins to organize into a stable integer mode. That integer might not be 6 at first. It might be 7 or 8 near onset, or it might be 5 if you are already driving strongly. The key is to recognize that you are in the right neighborhood if you see any stable integer mode.

Third, use small changes in ω to move the mode count. If you see 7–8 lobes and want 6, you are typically close. You want to shift the forcing so the m = 6 mode becomes the stable equilibrium. In many practical regimes, nudging ω upward or downward slightly can cause a mode transition. The direction depends on your specific geometry and damping, but the operating principle is consistent: make one small change, then wait long enough for the pattern to settle into its new attractor.

Fourth, once you land on m = 6, stop touching controls and verify persistence. This step matters more than getting the number. You want to see the hexagon remain stable for many rotations. You want the corners to be trackable. You want the pattern to show phase coherence rather than drifting into mode competition.

If the hexagon appears briefly and then collapses, treat that as valuable information. It means the system can access m = 6 but cannot remain there. The usual fixes involve reducing noise: improving lid damping, reducing ring speed ripple, smoothing the drive contact, or slightly increasing viscosity. It can also mean that you are sitting near a mode boundary. In that case, move slightly deeper into the m = 6 region by a small adjustment of ω.

Fifth, document your settings. Record Ω, ω, fluid depth, and ring geometry. The whole point of a regime diagram is reproducibility. A stable m = 6 state is a set of conditions you can return to. If you do not record them, you will end up "discovering" the hexagon repeatedly instead of learning the system.

A final tuning principle is patience. Mode locking is not instantaneous. The system needs time to exchange energy between wave modes and settle into a stable pattern. If you constantly adjust, you prevent that settling. The most effective tuning is often slow: change, wait, observe, record.

When you follow this protocol, the experiment feels less like guesswork and more like steering. You are not forcing corners into existence. You are guiding a jet through its instability landscape until it chooses the hexagon as its stable solution. That is the Saturn-like achievement: not the shape alone, but the controlled emergence of the shape from the dynamics.

10. Parameter Tuning for m = 6

If you can already produce an annular jet and you can already make it go polygonal sometimes, then getting a stable hexagon is rarely a matter of inventing new hardware. It is a matter of disciplined tuning. The tuning problem is not continuous; it is discrete. You are navigating between integer attractors in a system that wants to lock into one of several nearby modes. Your task is to place the system in a parameter region where m = 6 is the stable equilibrium—not just a transient that flashes past on the way to turbulence or to another polygon.

This section is deliberately hands-on. It treats Ω (background rotation) as the "planet setting," ω (differential ring rate) as the "wind forcing," and the remaining knobs (depth, viscosity, lid clearance) as "noise and sharpness controls." The practical logic is: set the planet, then tune the wind, then reduce noise only if you must.

10.1 A practical dialing guide: Ω first, then ω

The fastest route to m = 6 is to stop treating Ω and ω as interchangeable “speed knobs.” They do different jobs.

Ω is your regime selector. It determines whether the flow is disciplined enough to form a narrow jet and support a coherent wave. Too low, and you are in an everyday fluid regime where forcing tends to produce messy, quickly mixing motion. High enough, and you enter a regime where jets and low-mode waves can persist.

ω is your instability selector. It sets the strength of the shear and therefore whether the jet stays circular, develops a wave, and then locks into a polygonal mode. If ω is too small, nothing interesting happens. If ω is too large, you skip clean locking and go directly into irregular meanders and turbulence.

A practical dialing sequence looks like this:

First, choose an Ω that your platform can maintain smoothly with minimal wobble. In DIY practice, “highest possible” is not always best if it introduces wobble. A slightly lower but smoother Ω often produces a more stable polygon than a higher Ω with mechanical ripple. Your goal is a stable, quiet rotating frame.

Second, at that Ω, establish the jet at radius R with a small ω. Confirm that dye forms a narrow annular band and that the band stays narrow. If it does not, you are not ready to tune modes. Fix jet sharpness first (often by adjusting ring immersion or waiting longer for spin-up).

Third, increase ω in small increments until you see a coherent undulation. At this stage, stop thinking about m = 6. Your first target is “a clean integer mode at all.” That indicates you are in the polygon-supporting regime.

Fourth, once a stable integer mode appears, you adjust ω slightly to move between modes. This is where m = 6 becomes the target. You are now navigating a discrete landscape, and ω is the control that usually moves you across its boundaries.

Only after those steps do you touch other knobs. Depth, viscosity, and lid clearance can shift the mode boundaries and stabilize locking, but they are best treated as second-order controls. If you change too many things at once, you lose the map.

A practical mental image is: Ω sets the stiffness of the world; ω sets the stress on the jet. You want high stiffness and moderate stress. If stiffness is low, stress makes turbulence. If stress is low, stiffness yields a boring circle. The hexagon lives in the middle.

10.2 Reading the pattern: how to count lobes reliably

Counting lobes is deceptively tricky because the pattern is dynamic, not static, and because your eye is easily fooled by transient undulations. Reliable counting is essential, because you will make small tuning changes based on whether you are at m = 7 or m = 6.

The first rule is: count only after the pattern has stabilized. When a wave is still growing, it often shows multiple harmonics. You may see “extra bumps” riding on the main lobes. If you count then, you will mislabel the mode.

The second rule is: count crests, not corners. The polygon is a meander of the jet. Each lobe corresponds to a maximum outward or inward deflection of the dyed band. Corners appear as the jet path changes direction sharply, but those corners can be softened by diffusion or lighting. The more robust feature is the sequence of radial maxima or minima around the ring.

The third rule is: use video and slow motion. Top-down smartphone video allows you to pause at a moment when the wave looks clean and count without the distraction of motion. It also allows you to confirm that the same count persists over time.

The fourth rule is: measure angular spacing if you can. If the pattern is truly in mode m, then the lobes are roughly separated by 360 degrees divided by m. You do not need a protractor; even a simple reference mark and a rough count around the circle can confirm whether you are seeing a stable integer mode or a drifting mixture.

The fifth rule is: distinguish a traveling wave from a locked wave. A locked wave has a consistent structure that rotates at a steady rate relative to the rotating frame. If the wave is wandering, changing its lobe count, or showing intermittent splitting and merging, you are not in a stable attractor yet. Do not tune based on an unstable count.

A reliable sign you are in a stable mode is that the same lobe count remains visible over many full rotations of the pattern around the tank. If m changes spontaneously, you are in mode competition. If m stays constant, you have something you can tune.

10.3 If you get 7–8 lobes: how to move toward 6

Seeing 7 or 8 lobes is not failure; it is often the normal path toward 6. In many annular jet systems, higher mode numbers appear nearer the onset of instability, and lower mode numbers appear as forcing changes and the wave saturates more strongly. The exact direction depends on your geometry and damping, but the practical approach is consistent: move slowly, change one thing, and wait.

If you have a stable 7–8 lobe pattern and want 6, the first lever is ω.

Make a small adjustment of ω and then wait long enough for the system to settle. If the mode is stable, you will often see the pattern “snap” into a new lobe count rather than gradually morph. The snap is a signature of crossing a mode boundary.

If increasing ω makes the pattern noisier without changing m, you may be approaching turbulence rather than a mode transition. In that case, step back slightly and consider whether your jet is too broad or your drive too rough. A broad jet tends to support higher m and mode competition. A sharper jet tends to support cleaner low modes.

If decreasing ω causes the pattern to weaken and drift toward a nearly circular jet without transitioning to 6, you may be moving back below the instability threshold. That suggests you need to adjust in the opposite direction or modify damping so that the m = 6 mode becomes stable at lower amplitude.

Two secondary levers can help shift the system toward lower m (including 6):

Increase Ω slightly, if your mechanics allow it smoothly. Stronger background rotation can discipline the flow and sometimes favors lower mode numbers by sharpening jet structure and suppressing small-scale disturbances that support higher m.

Increase effective damping slightly, either by modestly increasing viscosity (a small glycerin fraction) or by reducing lid clearance to suppress surface-driven noise. Damping can discourage higher m modes that rely on finer structure and can help the system settle into a lower, more robust mode.

A useful diagnostic here is to look at jet width. If you see 8 lobes on a jet that is already broad and fuzzy, you will struggle to reach 6 by ω alone. The physical reason is that a broad shear zone supports multiple wavelengths. In that case, your first action is to sharpen the jet: reduce ring immersion depth, reduce ring width if possible, improve lid damping, and only then tune ω again.

In short: from 7–8 to 6, the primary action is small ω adjustments with patience; the secondary action is to sharpen and stabilize the jet so that lower modes have a clean waveguide to lock onto.

10.4 If you get 5 or fewer: how to move upward to 6

If you are seeing m = 5 or m = 4, you are still close, but you are on the other side of the mode landscape. Low m modes correspond to longer wavelengths around the ring, which often appear when forcing is strong and the pattern has saturated into large-amplitude meanders, or when the effective circumference is smaller than you think because the jet radius has drifted inward.

There are two common reasons you end up at 5 or fewer when you want 6.

The first is that ω is too high, pushing the system into a strongly nonlinear regime where the jet meanders with large amplitude and prefers lower m. In that case, reducing ω slightly can move you upward toward 6. Again, the key is to do it in small steps and wait. If you reduce ω too much, you may simply weaken the pattern without changing m.

The second reason is that your jet radius R has shifted. If the jet is not actually at the ring radius you intend, the circumference controlling mode selection may be different. A smaller effective R favors smaller m for the same wavelength scale. If dye shows the jet sitting noticeably inward from the ring's radius, you may need to adjust ring coupling or operating conditions so the jet is anchored where you want it.

To move from 5 to 6, the most direct strategy is to slightly reduce forcing or slightly increase the effective stiffness of the system:

Reduce ω slightly if you are clearly in a strong-meander regime.
Increase Ω slightly if you can do it smoothly; this can stabilize and sometimes raise mode number by sharpening the jet and reducing large-amplitude wandering.
Decrease viscosity if you have made the fluid too damped, because heavy damping can sometimes allow only the most robust large-scale mode to persist, which may be a lower m state.

The phrase "move upward to 6" should be interpreted carefully. You are not "adding a lobe." You are changing which wavelength the system finds energetically favorable. If your jet is too broad or too nonlinear, it will naturally settle into fewer, larger lobes. Your job is to pull it back into a regime where a slightly shorter circumferential wavelength is favored, allowing six lobes to fit stably.

A practical sign you are close to a 5-to-6 transition is when one of the lobes begins to split intermittently, or when the pattern shows a faint secondary undulation. That indicates mode competition. Small ω adjustments can then tip the system into stable m = 6.

10.5 Noise control: when to change depth, viscosity, or lid clearance

Depth, viscosity, and lid clearance are often called "secondary knobs," but they can be decisive when you are stuck near a mode boundary or when your system refuses to lock cleanly. The key is to change them only when the symptoms point to noise or smearing, not as a substitute for ω tuning.

Change lid clearance when surface effects are visible. If you see persistent surface ripples, shimmering reflections, or slosh-like oscillations that correlate with instability onset, your lid is not damping enough. Reducing lid clearance can suppress those waves and allow the jet

instability to saturate into a coherent polygon rather than into messy mode competition. If lid clearance is too small and the lid intermittently contacts the water, you will inject drag impulses and splashes, which is worse. The goal is close, consistent clearance.

Change viscosity when the instability breaks into small-scale turbulence too easily or when mechanical ripple seems to dominate. If you have a jittery polygon that never fully locks, a modest increase in viscosity can suppress small-scale eddies and smooth out the system's response. But if your jet is already broad and dye diffusion is already fast, adding viscosity may smear it further. Use viscosity changes as a stabilizer, not as a first response.

Change depth when your jet is either too fragile or too contaminated by boundary effects. If the flow seems dominated by bottom friction or if the jet cannot remain narrow, adjusting depth can change the balance between bulk flow and boundary layers. A slightly deeper layer can sometimes help by moving the active flow away from the bottom boundary and reducing the relative influence of bottom friction. A shallower layer can sometimes help by making the flow more quasi-two-dimensional, but it can also increase the relative importance of boundary drag. Depth changes are powerful but slow; they typically require draining and refilling, so use them after you have exhausted ω and minor lid adjustments.

A useful decision rule is:

If the mode count is stable but not the one you want, tune ω.
If the mode count is unstable and wandering, reduce noise first (lid clearance, mechanical smoothness, modest viscosity adjustment).
If the jet is not narrow and well-defined, fix jet formation (ring coupling, lid damping, Ω) before chasing m.
If you can access m = 6 but it collapses quickly, you are near the right regime but need better damping or sharper jet confinement.

The overall goal is to make the system "quiet enough" that it can choose its own integer. When the environment is noisy, the system does not select; it jitters. When the environment is disciplined, the system locks. That is what you are engineering: not a hexagon directly, but a quiet arena where the hexagon is the stable solution.

11. Diagnosing "Not Saturn-Like Yet" Outcomes

In a DIY build, the most discouraging moment is not "it failed." It is "it produced something interesting, but it isn't Saturn-like." That's the moment where the experiment can either become a learning loop or become a pile of half-remembered frustrations. The good news is that most "wrong" outcomes are not random. They are signatures. Each points to a specific

mismatch between what your apparatus is actually doing and what the Saturn-like target requires: a narrow annular shear layer at radius R that destabilizes into a low-mode traveling wave and then phase-locks into a persistent polygonal meander.

A Saturn-like hexagon is not merely any polygon in rotating water. It is a polygon that belongs to an annular jet. So the diagnostic question is always: where is the shear, what structure is the instability living on, and is the system quiet enough to lock?

This section is organized by the most common “not Saturn-like yet” outcomes and what they mean physically, followed by concrete fixes that preserve the run whenever possible.

11.1 Polygonal core vortex: why it happened and how to avoid it

A polygonal core vortex looks like a polygon, often spectacularly so, but it is not what you are targeting. It is a central vortex whose core boundary becomes unstable and takes on a polygonal shape. This can happen even in relatively small tanks and is often easier to achieve than an annular jet hexagon. That ease is exactly why it is such a common trap.

Why it happened is usually one of three reasons.

First, your forcing is effectively central rather than annular. If your differential driver is not truly localized at radius R—or if it couples strongly across the interior region—the resulting flow may concentrate vorticity at the center. The system then behaves like a vortex generator rather than a ring jet generator. The instability lives on the vortex core boundary, not on an annular shear layer.

Second, your jet never became narrow and anchored. If the shear you generate is broad and gradually varying with radius, it may naturally roll up into a single dominant vortex. Once that vortex dominates, the annular concept is lost. The “polygon” you see is a nonlinear vortex phenomenon, not a meandering jet.

Third, your ring coupling may be too deep or too aggressive, acting like a stirrer. Deep coupling injects angular momentum into a large volume of fluid quickly, which often produces a strong central vortex and secondary circulations. In that regime, the simplest stable structure is a vortex core.

How to avoid it is therefore also threefold.

Localize forcing at radius R. Ensure the ring actually imposes shear mainly at its radius and does not drag the entire interior. Gentle surface coupling is often better than deep coupling. If your ring dips too far into the water, reduce immersion depth.

Reduce the temptation toward vortex dominance by making the annulus visible and stable before you ever push into instability. If you can dye a stable annular band at radius R that stays

narrow and does not collapse inward, you are far less likely to end up with a core vortex outcome.

Increase the role of the rotating frame and decrease the role of “stirring.” This means ramping ω gently, keeping Ω smooth, and letting the system settle after each adjustment. Aggressive forcing tends to produce vortices. Controlled shear tends to produce jets.

A diagnostic that tells you whether you are in core vortex territory is where the strongest curvature and motion appear. In a core vortex polygon, the action is at the center and the polygon rotates like a rigid vortex boundary. In an annular jet polygon, the action is at radius R and the interior region can remain comparatively quiet. If the center is the show, you are not Saturn-like yet.

11.2 Axisymmetric shear with no polygon: likely causes and fixes

Sometimes you get the opposite frustration: the annular jet forms beautifully, the dye ring is narrow, everything looks disciplined, but the jet remains axisymmetric. No wave, no undulations, no polygon. This outcome is actually encouraging because it means your mechanical build and slosh suppression are doing their job. You are simply below the instability threshold or in a regime where the instability is suppressed.

There are several likely causes.

The most common is insufficient differential forcing. The jet exists, but the shear is not strong enough for perturbations to extract energy and grow. In this case, the fix is straightforward: increase ω in small steps, waiting between steps, until you see the first coherent undulation.

A second cause is excessive damping. If you are using a high-viscosity mixture, or if frictional boundary effects are strong, the disturbances that would become a wave may be damped before they can grow. This can also happen if the lid clearance is so tight that it introduces significant drag at the surface, effectively “braking” the jet.

A third cause is that your jet is too broad even though it looks clean. A broad shear zone can sometimes remain stable over a larger parameter range because it does not present a sharp waveguide for a single unstable mode. You get a smooth ring that resists forming a coherent low-mode wave. In this case, the solution is to sharpen the shear: narrower forcing ring, gentler but more localized coupling, or a slightly different ring immersion geometry.

A fourth cause is that your system is too perfect in an unhelpful sense: it lacks a perturbation to seed the instability. In many real systems, tiny disturbances are present naturally. In a very quiet experiment, it can take a long time for the instability to “find” a seed. A controlled way to seed is to introduce a very small asymmetry intentionally—such as a tiny pulse of dye injection at one

angle, or a brief, gentle change in ω—then return to steady forcing and let the system amplify what it wants to amplify. The key is to seed without stirring.

A practical diagnostic is to look for the earliest sign of instability: a faint standing or traveling undulation on the dyed ring. If you never see even a weak undulation after increasing ω gradually, you are either still below threshold or damping is dominating. Adjust ω first. If ω increases push you straight into turbulence without passing through a coherent wave, then your regime window is too narrow and you need more noise control or a sharper jet.

11.3 Turbulent breakdown: how to back out without losing the run

Turbulent breakdown is the “oh no” moment: the jet broadens, the pattern becomes chaotic, and your clean ring turns into a mottled mess. The instinct is to stop everything, drain, and start over. Often that is unnecessary. In many cases you can back out and recover a usable run if you treat turbulence as a temporary excursion caused by overforcing or a transient mechanical disturbance.

The first rule is not to chase turbulence with more forcing. Increasing ω rarely “forces the polygon through” turbulence. It usually makes turbulence worse.

The second rule is to reduce ω smoothly, not abruptly. Abrupt changes can generate additional transients. Step ω down to a calmer regime and keep Ω constant.

The third rule is to wait. Turbulence looks like the system is active, but the route back to coherence is often slow. The bulk needs time to re-establish a quasi-steady rotating frame and a narrow jet. If you immediately start re-tuning, you are adding new transients on top of the old ones.

If the turbulence was triggered by a temporary issue—like friction wheel slip, a small splash, or a motor hiccup—simply returning ω to its previous stable value and waiting can restore the jet. If the turbulence was triggered by crossing the true instability boundary into a chaotic regime, you may need to stay slightly below that boundary and approach it more gently.

A practical recovery sequence is:

Reduce ω to near-zero differential forcing while keeping Ω steady.
Wait until the flow looks calm and the dye motion becomes smooth again.
Re-establish the jet at low ω, confirm the dyed ring is narrow.
Increase ω again, but in smaller increments than before.

In other words, you do not lose the run just because you had turbulence. You lose the run if you panic and keep changing parameters rapidly. Patience and controlled back-out often salvage the session.

One important detail is that dye saturation can make recovery hard to interpret. If the tank has become uniformly dyed, you may not be able to see whether the jet is narrow again. In that case, switch to tracers, or inject a small fresh pulse of dye at radius R to re-mark the jet. The goal is to re-establish visibility, not to color the tank more.

11.4 Wobble and drift: mechanical sources versus fluid sources

Wobble and drift are especially confusing because they can masquerade as fluid dynamics. A slowly drifting polygon, a jet that seems to precess, or a lobe pattern that jitters can be caused by either fluid processes or by mechanical defects. The diagnostic skill here is to decide which it is, because the fixes are completely different.

Mechanical wobble has characteristic signatures.

It is phase-locked to the platform rotation. If you see a disturbance repeating at the same angular position in the lab frame every revolution, that is almost always mechanical. For example, one sector of the ring always looks slightly stronger, or the jet is always slightly closer to the wall at one angle. That indicates eccentricity or tilt.

It produces periodic modulation. If the polygon's amplitude strengthens and weakens rhythmically at a fixed frequency, suspect speed ripple in the drive or a physical wobble in the ring mount.

It correlates with audible or visible mechanical events. A friction wheel that squeaks intermittently or a belt that visibly oscillates often corresponds to a pattern jitter.

Fluid drift, by contrast, tends to be slower, less phase-locked, and more sensitive to parameter changes.

A traveling wave can drift relative to the rotating frame naturally. That drift is part of the phenomenon: the wave has a phase speed. If drift rate changes smoothly when you change ω, that is likely fluid. If drift persists regardless of ω and seems tied to a fixed lab-frame direction, that is likely mechanical.

Secondary circulations can cause the jet radius to drift inward or outward slowly. That is a fluid effect, often influenced by depth, viscosity, and boundary friction. It will not typically be pinned to a single angular position. Instead, the whole jet ring may slowly migrate.

Mode competition can produce lobe wandering and intermittent splitting. That is fluid. It often occurs near mode boundaries, and it often responds to small ω changes. Mechanical wobble tends not to produce true mode competition; it produces biased forcing.

A simple test is to mark a reference line on the lid or the lab frame and record video. If the "problem spot" always aligns with the same reference angle, it is mechanical. If it wanders, it is fluid.

Once you decide, you fix accordingly.

For mechanical wobble: re-level, re-center, improve ring concentricity, reduce runout, and improve motor smoothness. These are build fixes.

For fluid drift: adjust ω and Ω, sharpen the jet, change damping, adjust depth, or let the system settle longer. These are parameter fixes.

The reason this distinction matters is that many DIY builders try to tune away a mechanical defect by changing fluid parameters. That rarely works. You end up chasing your tail. Conversely, some builders rebuild hardware when the actual issue is that they are sitting near a mode boundary and need a small ω adjustment and more patience. The fastest progress comes from diagnosing correctly.

The unifying principle across all "not Saturn-like yet" outcomes is to ask what structure the instability is living on. If it is living on a core vortex, you built a vortex experiment. If it is living on a narrow annular jet, you built the right experiment and you are tuning the regime. If nothing is unstable, you are below threshold or overdamped. If everything is turbulent, you are above the clean-locking window or injecting mechanical noise. And if the pattern is biased to a fixed direction, the symmetry is broken mechanically.

Once you see outcomes as signatures rather than disappointments, each run becomes informative. That is exactly the mindset that turns a DIY Saturn-like hexagon from a curiosity into a repeatable laboratory analogue.

12. Data Collection and What Counts as Success

A DIY Saturn-like hexagonal jet becomes scientifically meaningful the moment you can reproduce it on demand and show that it is not an artifact of lighting, memory, or mechanical quirks. That requires data. Not "big data," but disciplined documentation: what the system was doing, how stable it was, what settings produced it, and how confidently you can claim it was m = 6 and not "six-ish."

The point of data collection in this project is not to impersonate a professional lab. It is to make your DIY work falsifiable and repeatable. If another builder can match your geometry and your dimensionless regime knobs and get the same locked mode, you have done real laboratory analog science.

This section defines success in operational terms and describes the smallest set of measurements that let you support that claim.

12.1 Documenting stability: duration, rotation rate, and phase-locking

“Success” is not a snapshot of a hexagon. Success is a stable, phase-coherent m = 6 annular jet pattern that persists over time under steady operating conditions.

The first success metric is duration. You should define duration in a way that is meaningful for rotating systems: count rotations or count pattern cycles rather than measuring only wall-clock time. A hexagon that persists for 30 seconds at high pattern speed may represent many pattern cycles. A hexagon that persists for 30 seconds but drifts wildly in shape may not be stable. Thus, the best stability record includes both time and number of rotations.

A practical stability record includes:

The start and end time of the stable m = 6 interval.
The number of complete pattern revolutions during that interval (if the pattern rotates relative to the lid).
Whether the lobe count remained constant (no splitting/merging events).
Whether the pattern’s amplitude stayed roughly constant or grew/decayed.

The second success metric is rotation rate. Record the platform rotation rate Ω and the differential ring rotation rate ω as directly as you can. In a DIY setting, you may measure in rpm. That is fine. What matters is consistency and clear documentation of how you measured it. If your platform has a dial with no numerical readout, add your own measurement marks and use an optical tachometer or video-based counting.

The third success metric is phase-locking. In the Saturn analogy, “persistence” is the story: the hexagon remains a coherent structure rather than drifting into chaos. In a tabletop system, phase-locking can mean different things depending on your frame of reference:

Locked in the rotating frame: the hexagon is stationary relative to the co-rotating lid. This is the cleanest demonstration of a locked wave state.
Locked drift: the hexagon rotates at a steady angular speed relative to the lid, maintaining its shape while traveling. This is still Saturn-like if the pattern is coherent and stable.
Not locked: the pattern’s phase speed fluctuates, corners wander, and lobes appear/disappear. This is not a success state.

How to document phase-locking cheaply is to include a reference mark in your video. Put a visible mark on the lid or tank rim and record the pattern from above. Then, in playback, track whether a specific corner stays at a fixed angle, drifts steadily, or jitters. If it drifts steadily, you

can quantify the phase speed by measuring how long it takes a corner to pass the reference mark.

A strong “success documentation clip” is a continuous top-down video segment showing at least several full cycles with an obvious reference mark, plus separate short clips of the tachometer readings for Ω and ω.

12.2 Quantifying the jet: radius, width, and wave amplitude

If you want your results to be reproducible, you must describe the jet as a physical object, not just as “a ring.”

There are three core measurements: jet radius R, jet width W, and wave amplitude A.

Jet radius R is the mean radius of the dyed annular band when the jet is established. This is the fundamental scale for mode selection because circumference is $2\pi R$. Measure R by placing a scale reference in the image plane (for example, a ruler taped under a transparent tank or marks on the lid) and then extracting radius in pixels from video frames. If you do not want to do pixel measurement, you can do a direct physical measurement by marking the ring radius on the lid and measuring from center, but image-based measurement is usually more reliable and yields a record.

Jet width W is the thickness of the dyed band. A narrower jet generally corresponds to cleaner mode selection and more stable polygons. In image terms, W is the radial thickness of the high-contrast region. You do not need to resolve it to a tenth of a millimeter. You need a consistent definition, such as “full-width at half-maximum intensity” if you are doing image analysis, or “visually estimated thickness between the inner and outer steep dye gradients” if you are doing a simple approach. The key is consistency across runs.

Wave amplitude A is the radial deviation of the jet from its mean radius. For a polygonal jet, you can define amplitude as half the difference between the maximum and minimum radius of the jet path at a given time. This is especially useful because it distinguishes a barely-detectable undulation from a saturated polygon. It also lets you track whether the hexagon is stable or slowly changing in strength.

In a DIY workflow, the simplest quantitative method is to select a few representative frames from the stable interval and measure:

R_{mean}: average radius around the circle.
R_{max} and R_{min}: maximum and minimum radial positions of the dyed band.
$A = (R_{max} - R_{min}) / 2$.
W: approximate width of the band near the midpoints between lobes, where curvature is lower and measurement is easier.

You can do this with basic image tools or even by importing frames into a simple drawing app that lets you measure pixel distances. What matters is that you store the frames and your measurement method, not that you reach lab-grade precision.

Also record fluid depth H and ring geometry (ring radius and ring width), because these define the physical forcing and confinement.

The more convincing your “Saturn-like” claim becomes, the more it will hinge on showing that you had an annular jet of known radius and width, not just “a polygon somewhere in the flow.”

12.3 Simple uncertainty handling for DIY instrumentation

Uncertainty handling in DIY science should be humble and explicit. The goal is not to impress; the goal is to prevent self-deception. You want to know whether a difference you see between runs is real or within measurement noise.

A simple uncertainty protocol has three parts: repeat measurement, estimate resolution, and propagate only what matters.

Repeat measurement means you measure the same quantity multiple times, ideally across multiple frames and multiple runs. For example, measure Ω three times during a run and confirm it is stable. Measure R_mean on three frames during the stable interval and confirm it does not drift significantly. If it drifts, that drift is part of your data.

Estimate resolution means you state your measurement precision. For a tachometer, your resolution might be ±0.1 rpm or whatever your device supports. For image-based radius measurement, your resolution might be “±2 pixels,” which translates to a physical distance depending on your scale reference. The important thing is to record it.

Propagate only what matters means you do not need a full formal error analysis. You need to be honest about which comparisons are meaningful. For example:

If your measured R differs by less than your radius resolution between runs, treat the difference as negligible.
If your measured Ω varies by less than the tachometer resolution, treat Ω as constant for that run.
If your lobe count changes, that is not an uncertainty problem; it is a regime change.

A particularly useful DIY uncertainty practice is to create a run log sheet where you record: Ω, ω, H, viscosity mix, lid clearance setting, ring immersion depth, observed mode, stability duration, and notes on mechanical behavior. The act of writing it down reduces the tendency to “remember” a run as better than it was.

A final note: the lobe count m is often the most robust variable because it is discrete. But even here you can have uncertainty if the pattern is unstable. So your uncertainty handling includes a categorical statement: “m = 6 stable” versus “m fluctuating between 6 and 7.” That distinction is itself a form of uncertainty management.

12.4 A “minimum publishable dataset” for reproducible DIY science

The minimum publishable dataset is the smallest set of information that allows another competent DIY builder to replicate your result and judge whether they achieved the same phenomenon.

It is not a single photo. It is a package.

A minimal dataset includes:

A clear description of geometry: tank radius, jet radius target, ring width, ring immersion depth, lid clearance, fluid depth.
Fluid composition and temperature: water or water–glycerin ratio, approximate temperature, any surfactants or additives.
Rotation settings: Ω and ω with measurement method and typical variation during the run.
A top-down video clip of the stable m = 6 interval with a reference mark visible, at known frame rate if possible.
Representative still frames from the stable interval used for measurement.
Measured jet properties: R_mean, W, and A for at least a few frames, plus how you measured them.
Stability metrics: duration in time, number of pattern cycles, and whether phase speed was steady.
A run log for at least several runs showing how you reached the stable hexagon, including near-miss runs (m = 7, m = 5, turbulence) because those define the neighborhood of the regime.

If you include that package, you have transformed a hobby experiment into a reproducible laboratory analogue. Another builder does not need your exact motor brand to replicate. They need your geometry, your operating regime, and your disciplined documentation.

What counts as success under this standard is:

A clearly annular jet (not a core vortex).
A stable m = 6 lobe count for a meaningful interval.
A coherent, phase-consistent pattern (locked or steadily traveling).
Measured and reported Ω, ω, and jet geometry sufficient to reproduce.

That is the essence of DIY science: not a spectacle, but a repeatable phenomenon with enough structure that someone else can test it. When you reach that point, you are not only making a Saturn-like hexagon—you are making a tabletop system that earns the word “laboratory.”

13. Interpretation: What This Teaches About Planetary Atmospheres

A DIY Saturn-like annular jet hexagon is not a toy because it looks like Saturn. It is a toy because it isolates a small set of mechanisms—rotation, shear, wave selection, and nonlinear saturation—so cleanly that you can watch a planetary-scale logic unfold on a tabletop. The interpretation value is therefore not “we recreated Saturn.” It is: we recreated a dynamical *type* that Saturn appears to instantiate.

In planetary atmospheres, almost everything matters: stratification, radiation, microphysics, deep convection, topography (for rocky planets), seasonal cycles, chemistry, and the fact that a “surface” might not exist. In your lab, almost none of that exists. What you have is a controlled rotating fluid with an annular shear layer and the ability to tune forcing. This is both the limitation and the power. It is a stripped-down model that answers a very specific question: under what conditions does a jet prefer a small-integer polygonal wave and remain coherent over long times?

This section interprets what that means and how to think about the Saturn hexagon through the lens of your setup.

13.1 What the analogue captures well (and what it cannot)

What the analogue captures well is the kinematics and dynamics of a rotating shear layer selecting a preferred circumferential wave and saturating into a stable, coherent pattern.

It captures quasi-two-dimensionality under strong rotation. In both the lab and in a rapidly rotating atmosphere, rotation suppresses some types of vertical motion and can make flows organize into large-scale horizontal structures—jets, vortices, and waves—rather than immediate three-dimensional turbulence.

It captures annular jets as waveguides. A narrow jet is not just a region of fast flow. It is a structure with gradients that can support and amplify particular wave modes. The jet is where the wave “lives.” This is the essential Saturn-like ingredient: the hexagon is not a solid object; it is a wave pattern embedded in a jet.

It captures mode selection and discrete wavenumbers. When you tune ω and watch the system jump between m = 7 and m = 6, you are seeing the same qualitative fact Saturn’s hexagon suggests: there is a dynamical reason small integers can be preferred. The atmosphere does not

choose an arbitrary number of corners. It chooses a stable wave number given the jet geometry and forcing.

It captures persistence as a property of a locked nonlinear state, not as a property of rigidity. Your hexagon persists because the wave has saturated into a stable balance: forcing feeds it, damping bleeds it, and the resulting pattern is an attractor. That is exactly the kind of persistence that can occur in a planetary atmosphere without any "solid" structure.

What the analogue cannot capture is equally important.

It cannot capture Saturn's vertical stratification and thermal forcing profile. Saturn's atmosphere is layered, with varying stability, radiative timescales, and latent heat effects. Your tank is at best weakly stratified unless you deliberately add thermal gradients or density layering.

It cannot capture compressibility and realistic atmospheric equations of state. Saturn's atmosphere is a compressible gas; your fluid is effectively incompressible.

It cannot capture the full planetary geometry and the polar environment. Saturn's hexagon is at high latitude, where the Coriolis parameter and effective geometry vary with latitude. Your tank is a flat plane with a single rotation rate. The "north pole" in your experiment is a symmetry axis, not a spherical pole with curvature and latitude-dependent effects.

It cannot capture Saturn's long-term energy balance. Planetary persistence involves continuous energy sources (radiation, convection) and sinks over long times. Your forcing is motor-driven and externally controlled.

So the correct interpretive stance is: your analogue is a mechanism demonstrator. It is excellent for testing "can rotation + annular jet + nonlinear wave locking yield a persistent hexagon?" It is not designed to answer "what exact microphysical or radiative processes sustain Saturn's jet?" Both are valid questions; your experiment speaks strongly to the first.

13.2 Persistence, forcing, and the meaning of a selected wavenumber

In your experiment, persistence is a relationship between energy injection and dissipation under a geometric constraint.

Energy is injected by differential rotation: the ring imposes shear that creates and maintains the jet. Energy is dissipated by viscosity, boundary friction, and any small-scale turbulence that emerges. The pattern persists when the wave mode that forms can draw energy from the shear efficiently while remaining stable against breakup.

The selected wavenumber m is therefore not an aesthetic choice. It is the system's answer to: which circumferential wavelength fits the jet and can be maintained by the available forcing without destabilizing into chaos?

The most important conceptual shift here is to stop thinking of m as “number of corners” and start thinking of m as “the wave count around the circumference that the jet supports.” The jet’s radius sets the circumference. The jet’s width and shear profile set which wavelengths are unstable and which saturate cleanly. Forcing and damping set how far into nonlinearity the system goes. When you change ω, you change the balance, and the preferred m can shift.

This is why the pattern changes in discrete jumps. Each m corresponds to a different wave configuration that can be stable in a certain parameter band. You are not tuning a continuous shape. You are navigating between discrete attractors.

In planetary terms, this reframes “persistence” as a dynamical equilibrium rather than a mysterious rigidity. A persistent hexagon is a wave attractor of a jet in a rotating system with stable forcing. If the forcing and damping remain roughly steady, the attractor can persist for long times.

This also suggests a more subtle interpretation: “persistent” does not mean “unchanging.” It means “self-restoring.” Your tabletop hexagon may wobble slightly, corners may breathe, amplitude may fluctuate, but the overall m = 6 structure returns because it is the stable mode. That is a hallmark of a nonlinear attractor maintained by continuous forcing.

13.3 Why a north-pole hexagon can be stable without being rigid

The Saturn hexagon is often described in language that tempts rigid interpretation: a giant hexagon “sitting” at the pole. But your lab experience should inoculate you against that picture.

In the lab, the hexagon is not a material object. It is a pattern in a moving fluid. It can be stable for long times while the fluid composing it is continuously replaced. A corner is not a piece of fluid that stays there. It is a location in the pattern where the jet’s trajectory bends in a consistent way. The fluid parcels pass through; the pattern remains.

This is the deepest lesson your analogue teaches: stability is not rigidity. Stability is invariance of form under flow.

A hexagon can be stable if the underlying jet and wave satisfy a balance: the wave’s phase speed, the jet’s shear, and the nonlinear saturation cooperate so that the wave neither damps away nor steepens into turbulence. That balance can hold while individual fluid elements move through the pattern. In other words, the hexagon is a standing or steadily traveling wave locked to the jet, not a rigid “structure.”

The “north pole” aspect adds an additional stabilizing intuition. Polar regions in rotating systems can support strong, coherent vortices and jets because the geometry and rotation constrain motion. In your tank, the axis of rotation plays a similar organizing role: it is a symmetry anchor.

When you build a jet at radius R and the system is quiet, the symmetry anchor helps the wave lock into a coherent mode.

A helpful mental model is: the hexagon is a wave that uses the jet as a track and uses rotation as a stabilizing constraint. It persists not because it is hard, but because it is dynamically preferred and continuously maintained.

Your experiment also teaches why "small wobble does not contradict stability." In the lab, even a stable m = 6 mode can breathe slightly. That breathing is not failure; it is the natural response of a nonlinear system under mild noise. The stability claim is about the persistence of the mode number and the coherent structure, not about geometric perfection.

That is likely the right way to interpret Saturn's hexagon as well: stable mode selection does not require a perfect rigid hexagon; it requires a stable dynamical state that keeps returning to the same mode.

13.4 How the lab setup reframes the question "why a hexagon?"

"Why a hexagon?" is often asked as if the hexagon itself is the primary mystery. Your lab setup reframes it into a better question: "Why does this jet select m = 6 in this regime?"

That reframing is powerful because it replaces a single-object puzzle with a regime-based explanation. Instead of hunting for a one-off causal story, you look for:

What defines the jet's effective circumference and shear profile?
What instability family is available in this flow?
Which modes are unstable and which saturate cleanly?
What forcing and damping balance supports long-lived wave locking?
What parameter changes would shift the preferred m?

In your tank, you can answer these questions by tuning ω and observing mode transitions. You can show that m changes with forcing. You can show that there is a band where m = 6 is stable. You can show that m = 6 is not "special" in isolation; it is special because your jet geometry and forcing place it in the stable window.

That insight carries directly to planetary interpretation: the hexagon is not a magical number. It is a stable wave number for a particular jet under a particular set of planetary conditions.

The lab also suggests a more nuanced perspective on uniqueness. Saturn's hexagon is famous partly because it is visually striking and persistent. But your lab demonstrates that polygonal jet patterns can be generic in rotating shear flows. The "mystery" is not that polygons exist; the mystery is why Saturn's polar jet sits in a regime that keeps it locked, and why it sits specifically in the m = 6 attractor rather than wandering among nearby modes.

In that sense, the lab analogue converts a sensational question into a scientific one: not “why does Saturn have a hexagon?” but “what maintains a polar jet in a stable low-mode wave state, and what fixes the selected mode number over time?”

Your DIY build cannot answer every part of that planetary question, but it does something more fundamental: it demonstrates that a stable, persistent hexagon can emerge from simple rotational fluid dynamics without requiring rigid structures or exotic explanations. That alone is a major interpretive lesson—and it is exactly the kind of lesson laboratory analogues are meant to teach.

14. Extensions and Variants

Once you can reliably produce a Saturn-like annular jet polygon—especially m = 6—you have crossed the hardest threshold: you now have an instrument, not a one-off demo. Extensions are valuable for two reasons. First, they let you probe “why” more deeply by changing the shear profile, stratification, or boundary conditions. Second, they let you make the experiment more robust, more visually persuasive, or more portable without losing the mechanism.

A guiding principle for extensions is to preserve the core dynamical ingredients: a strongly rotating background, a narrow annular shear layer, and a parameter window where a circumferential wave can saturate and phase-lock. Everything else you add should either make the jet more Saturn-like or make the lab state easier to see and reproduce.

14.1 Two-ring systems and more Saturn-like shear profiles

A single ring drive is the minimal annular jet generator, but it produces a shear profile that is often “too simple” compared to planetary jets. Real atmospheric jets can have multiple shear layers: a fast core with slower flow on both sides, sometimes with adjacent counterflow or secondary jets. A two-ring system is the simplest DIY way to approximate that complexity.

The idea is to use two concentric rings at different radii, each capable of a controlled differential rate relative to the co-rotating tank. There are several ways to configure them:

Inner ring faster, outer ring slower: This creates a shear layer between them and can sharpen a jet at an intermediate radius.
Outer ring faster, inner ring slower: This can create a jet nearer the outer ring or invert the shear gradient, sometimes changing which modes are favored.
Counter-rotating rings: This creates a stronger shear region and can make the jet extremely sharp, but it also increases the risk of turbulence if applied aggressively.

The most Saturn-like use of two rings is not necessarily to drive them wildly differently, but to shape the velocity profile so you have a pronounced jet core with gentler gradients outside it. This often yields cleaner mode selection because the wave has a better-defined waveguide.

Mechanically, two rings are a step up in complexity because concentricity and wobble control become more demanding. A practical approach is to keep one ring as the primary driver and use the second as a "trim" ring that makes small adjustments to the surrounding flow. If you treat the second ring as a gentle shaper rather than a second aggressive stirrer, you can gain realism without sacrificing stability.

From an interpretation standpoint, two rings let you ask a deeper question: is m = 6 robust to changes in shear profile, or is it a fragile feature of one specific forcing geometry? If m = 6 persists across multiple shear profiles, you have a stronger analogue claim: the selected wave number is tied to the jet geometry and rotation regime, not to a particular mechanical trick.

14.2 Bottom topography or thermal forcing as additional realism

Planetary atmospheres are not uniform layers. They are stratified and forced by heat. Your base experiment is typically a single-layer, near-isothermal fluid. You can add realism in two main ways: by shaping the bottom boundary (topography) or by imposing thermal forcing (heating/cooling gradients).

Bottom topography is the simpler of the two. Adding a gentle bump, ridge, or sloped bottom changes the flow's effective depth and can introduce preferred radii or barriers to radial transport. In a rotating system, depth variations can influence vorticity behavior and can stabilize or destabilize jets. Even a simple conical insert can change how the interior region communicates with the jet region.

The danger with topography is that it can break symmetry in ways that obscure your mechanism. The Saturn-like target is a wave on an annular jet, not a wave pinned to a bump. So if you add topography, it should be smooth, axisymmetric, and centered. The goal is to add realism without adding an artificial "corner maker."

Thermal forcing is more Saturn-like in spirit because planetary jets are ultimately tied to energy flows. In a DIY context, thermal forcing can be implemented as heating at one boundary and cooling at another, or by placing a cold pack beneath part of the tank while gently heating elsewhere. But thermal forcing is also more complex because it introduces buoyancy-driven circulations that can compete with the quasi-two-dimensional rotation regime.

The safest thermal extension is to keep it mild and axisymmetric: for example, slightly warming the outer region or slightly cooling the center in a symmetric way, then observing whether the

annular jet's stability changes. The goal is not to create vigorous convection, but to introduce a slow, controlled stratification or baroclinic tendency.

These extensions let you explore a powerful interpretive question: is the hexagon fundamentally a shear-wave phenomenon that persists under added realism, or does added realism destabilize it? Either outcome teaches something. If it persists, it suggests robustness. If it breaks, it suggests Saturn's environment may be doing subtle stabilizing work that your minimal model lacks.

14.3 Scaling up: when a larger tank dramatically improves clarity

If there is one extension that almost always helps, it is scaling up. A larger tank improves the experiment's clarity disproportionately because it separates scales.

In small tanks, everything is close: boundaries, bottom friction, surface effects, and mechanical imperfections all interact on similar length scales. The jet's circumference is small, so small geometric errors become large fractional errors. The pattern has less room to develop before it runs into boundary influences.

In a larger tank, the jet can be narrower relative to the overall size, and the polygonal wave can have more "clean" circumference. Boundary layers occupy a smaller fraction of the domain. Small mechanical wobble becomes a smaller fractional perturbation. Visualization improves because you have more pixels per feature in top-down video.

Scaling up also improves the stability window. In many DIY builds, the problem is not that m = 6 is impossible. It is that the stable m = 6 window is narrow and easily disrupted. A larger system often broadens that window because the wave has more space to organize and because the system's effective noise floor is lower.

There are practical constraints: a larger tank means more mass, more structural rigidity needed, and more safety considerations for rotating inertia. But if your goal is clarity and repeatability, scaling up is one of the highest-impact investments you can make once you have proven the concept at smaller scale.

A particularly effective intermediate step is not just "bigger," but "bigger and stiffer." A larger tank on a rigid, well-balanced platform can make the difference between "sometimes a hexagon" and "hexagon on demand."

14.4 Classroom and outreach versions that still preserve the physics

A classroom or outreach version is successful if it preserves the mechanism while reducing build complexity and safety risks. The temptation is to oversimplify until the hexagon becomes a fragile lucky event. The better approach is to keep the key physics intact but adjust

expectations: you may not always get perfect m = 6, but you should reliably get annular jets and polygonal modes, and you should be able to explain mode selection visually.

A good outreach version has four features:

Low voltage, battery-powered, and guarded. Safety must be intrinsic, not dependent on perfect behavior.
Clear top-down viewing. A transparent lid and a fixed camera or projector feed helps a group see the pattern.
Fast spin-up and short run cycles. Classroom time is limited; the system should reach a stable jet quickly.
Robust visuals. Fluorescein under blue light or strong dye contrast makes the effect visible from a distance.

To preserve the physics, you must preserve controlled differential rotation. Many outreach attempts fail by using stirring paddles or by rotating only the container without a clear annular forcing mechanism. That produces vortices and swirls, but not Saturn-like annular jet polygons.

A workable simplification is to use a single, robust friction wheel drive for the ring, with conservative ω settings and a strong lid to suppress slosh. The demonstration goal can be: show an annular jet, show a transition to a polygonal mode, show mode switching with small ω changes, and explain why integers appear. If m = 6 appears, it is a highlight, not the sole success criterion.

If you design the outreach version as an “integer mode selection machine” rather than as a “hexagon machine,” you preserve the scientific heart. You can then connect that heart back to Saturn: the hexagon is one selected mode among several, and its persistence is an attractor phenomenon, not a frozen shape.

These extensions—two-ring profiles, added realism through boundary and thermal modifications, scaling up, and outreach adaptation—are not optional flourishes. They are how you turn a DIY build into a miniature research platform: a device that can answer progressively deeper questions about why jets select polygons, how those polygons persist, and what that teaches us about the strange stability of patterns in rotating worlds.

15. Conclusion: Building a Persistent Hexagon as Hands-On GFD

The point of this project was never “make a cool shape in a tank.” The point was to build a tabletop system that behaves like a planetary atmosphere in one very specific and illuminating way: a narrow annular jet in a rapidly rotating frame selects a small-integer circumferential wave and then holds it as a persistent, coherent pattern. When you achieve a stable m = 6

annular-jet hexagon, you have not copied Saturn's north pole. You have reproduced a *mechanism class*—rotation-imposed quasi-2D flow, localized shear, instability with discrete mode selection, and nonlinear saturation into a locked wave state. That is the essence of hands-on geophysical fluid dynamics: a simple apparatus that turns an abstract explanation into something you can see, tune, and measure. And once you can do it repeatably, you have an instrument that can ask real questions—how modes switch, what stabilizes persistence, and what it means for a hexagon to be stable without being rigid.

15.1 One-paragraph synthesis

A Saturn-like hexagon in the lab is a phase-locked wave living on a narrow annular jet inside a strongly rotating fluid. The operational recipe is: first build a quiet rotating world (smooth Ω, minimal slosh, good symmetry), then create a localized shear layer at radius R (a truly annular driver, not a stirrer), then let a coherent circumferential wave grow and saturate without tipping into turbulence (slow ω ramps, patience between adjustments), and finally stabilize the chosen mode with noise control rather than brute forcing (lid clearance, jet sharpness, modest viscosity or depth tweaks only when needed). In that regime, the pattern persists because forcing and dissipation balance in a stable attractor: fluid parcels move through it, but the wave form remains. What you learn is not merely that polygons can appear, but that small integers can be *selected* by a jet as a dynamical solution—and that "persistence" is the signature of locking, not of rigidity.

15.2 Practical takeaways: the shortest path to a stable m = 6

If the goal is m = 6 on demand, the shortest path is not "more power" or "more complexity." It is discipline and sequencing. The takeaways below are the ones that most directly compress the learning curve.

First, prioritize smoothness over speed. A slightly lower Ω that is clean and wobble-free will beat a higher Ω that injects periodic forcing. The hexagon is an attractor in a quiet system; noise pushes you into mode competition or turbulence.

Second, do not tune polygons until the jet is a narrow, visible annulus at radius R. If you cannot dye a clean ring that stays narrow for multiple rotations, you are not yet tuning the right phenomenon. Fix the jet first: ring coupling depth, ring width, lid damping, and patient spin-up.

Third, ramp ω like a scientist, not like a driver. Small increments, long waits. The system responds on timescales that are long compared to knob-twisting. If you change ω repeatedly before the pattern settles, you will overshoot the stable window and blame the physics for what was actually impatience.

Fourth, treat m as a regime indicator. Seeing 7–8 lobes means you are close; it is not failure. Seeing 5 means you are close from the other side. The "map" is built by logging: which ω produced which m at a fixed Ω. Once you have that map, m = 6 becomes a destination you can return to, not a lucky accident.

Fifth, use noise control only when the symptom is noise. If m is stable but not 6, adjust ω. If m wanders, corners jitter, or the jet breaks intermittently, reduce noise: tighten lid clearance (without contact), improve drive smoothness, consider a modest viscosity increase, or adjust depth. Noise control is how you hold the attractor once you find it.

Sixth, define success as persistence plus phase coherence, not geometric perfection. A stable hexagon that breathes slightly but holds m = 6 and a consistent phase speed is more Saturn-like than a sharp-looking snapshot that collapses a few seconds later.

If you follow those rules, the shortest path is usually: water first, strong lid damping, establish clean jet, find any stable polygonal mode, then walk ω gently until the system snaps into 6, then stop touching controls and let it prove its stability.

15.3 What to explore next: controlled experiments you can run at home

Once you have a stable m = 6 state, you can do real "home GFD" experiments by changing one variable at a time and watching what the system selects. The projects below are controlled, incremental, and genuinely informative.

Mode boundary mapping. At a fixed Ω and fixed geometry, sweep ω slowly across the polygonal regime and record where m transitions occur (8→7→6→5, etc.). Your output is a regime diagram: ω ranges that support each stable m and where mode competition appears.

Hysteresis tests. Approach m = 6 from above (starting at m = 7–8) and from below (starting at m = 5) and see whether the transition ω is the same in both directions. If it differs, you have demonstrated hysteresis—classic nonlinear attractor behavior.

Persistence under controlled noise. Introduce a small, repeatable perturbation—like a tiny step change in ω or a brief pulse of dye injection—and measure whether the hexagon returns to m = 6. This tests whether your state is truly a restoring attractor.

Jet sharpness experiments. Change only one thing that affects jet width (ring immersion depth, ring contact pressure, or viscosity) and measure how W correlates with mode stability and with the width of the m = 6 parameter window. This connects "how sharp is the jet?" to "how clean is mode selection?"

Phase-speed measurement. With a reference mark, measure how fast the hexagon rotates relative to the lid as you vary ω slightly within the m = 6 window. You will learn whether you have a nearly stationary locked state or a traveling wave and how sensitive it is to forcing.

Two-color shear visualization. Inject two dyes, one just inside R and one just outside, and watch how the interface behaves in different modes. This makes the shear layer visible and helps distinguish true jet dynamics from bulk swirling.

Scale and realism variants. If you can safely scale up, repeat the same parameter sweep in a larger tank and compare: does the m = 6 window widen, does persistence increase, do transitions sharpen? Or, add an axisymmetric bottom slope and see whether mode selection shifts without introducing artificial pinning.

Every one of these can be done at home with the same apparatus. Each produces a dataset that is small but meaningful: not only “I made a hexagon,” but “here is when it appears, how it persists, how it switches, and how robust it is.” That is the real payoff. You end up with a hands-on understanding of why planetary atmospheres can stabilize strange, beautiful structures—and you can demonstrate that understanding with a system you built yourself.

16. References

16.1 Planetary observations and Saturn hexagon overview sources

- Godfrey, D. A. (1988). *A hexagonal feature around Saturn's north pole*. Icarus, 76(2), 335–356. https://doi.org/10.1016/0019-1035(88)90075-9.
- Allison, M., Godfrey, D. A., & Beebe, R. F. (1990). *A wave dynamical interpretation of Saturn's polar hexagon*. Science, 247(4946), 1061–1063. https://doi.org/10.1126/science.247.4946.1061.
- Baines, K. H., et al. (2009). *Saturn's north polar cyclone and hexagon at depth revealed by Cassini/VIMS*. Planetary and Space Science, 57(14–15), 1671–1681. https://doi.org/10.1016/j.pss.2009.06.026.
- Sánchez-Lavega, A., et al. (2014). *The long-term steady motion of Saturn's hexagon and the stability of its enclosed jet stream under seasonal changes*. Geophysical Research Letters, 41, 1425–1431. https://doi.org/10.1002/2013GL059078.
- Fletcher, L. N., et al. (2018). *A hexagon in Saturn's northern stratosphere surrounding the emerging polar vortex*. Nature Communications, 9, 3564. https://doi.org/10.1038/s41467-018-06017-3.
- Sánchez-Lavega, A., et al. (2020). *An enduring rapidly rotating hexagon-shaped wave around Saturn's north pole*. Nature Communications, 11, 1610. https://doi.org/10.1038/s41467-020-16110-1.
- Pryor, W. R., et al. (2019). *Detection of Saturn's north polar hexagon in Cassini UVIS auroral observations*. Journal of Geophysical Research: Planets, 124, 1979–1988. https://doi.org/10.1029/2019JE005922.
- Sanz-Requena, J. F., et al. (2018). *Dynamics of Saturn's polar regions from Cassini imaging: jets, vortices, and long-lived waves*. Icarus, 304, 27–49. https://doi.org/10.1016/j.icarus.2017.12.025.
- Antuñano, A., et al. (2018). *Cloud morphology and dynamics within Saturn's hexagon from Cassini ISS*. Icarus, 299, 264–279. https://doi.org/10.1016/j.icarus.2017.07.010.
- Antuñano, A., et al. (2019). *Potential vorticity structure and wave activity in Saturn's polar regions*. Journal of Geophysical Research: Planets, 124, 363–384. https://doi.org/10.1029/2018JE005711.
- Baines, K. H., et al. (2018). *The eye of Saturn's North Polar Vortex*. Geophysical Research Letters, 45, 11,022–11,028. https://doi.org/10.1029/2018GL078020.

- Ingersoll, A. P., et al. (2020). *Cassini exploration of the planet Saturn*. Space Science Reviews, 216, 122. https://doi.org/10.1007/s11214-020-00751-1.
- Hueso, R., et al. (2020). *Saturn atmospheric dynamics one year after Cassini: new constraints and open questions*. Icarus.
- Wong, M. H., et al. (2021). *Interaction of Saturn's Hexagon with a series of convective storms in 2014*. Geophysical Research Letters.
- Sayanagi, K. M. (2018). *Saturn's polar atmosphere*. In A. M. King, et al. (Eds.), *Saturn in the 21st Century*. Cambridge University Press.
- NASA (2024). *Saturn's "Hexagon in Motion" (Cassini imagery feature/update)*.
- ESA / NASA / JPL-Caltech / SSI (2018). *Saturn's famous hexagon may tower above the clouds* (Cassini result highlight).
- ESA (2012). *Saturn's hexagon: movie/visualization* (Cassini-based animation).
- The Planetary Society (2010). *Saturn's hexagon re-created in the lab* (accessible overview).

16.2 Rotating-tank annular jet and polygonal wave experiments

- Barbosa Aguiar, A. C., Read, P. L., Wordsworth, R., Salter, T., & Yamazaki, Y. H. (2010). *A laboratory model of Saturn's North Polar Hexagon*. Icarus, 206(2), 755–763. https://doi.org/10.1016/j.icarus.2009.10.022.
- Morales-Juberías, R., et al. (2011). *Saturn's hexagon: a rotating-tank analogue and jet-wave interpretation* (Cassini/lab-linked analysis). Icarus.
- Morales-Juberías, R., Sayanagi, K. M., Simon, A. A., & Cosentino, R. G. (2015). *Meandering shallow atmospheric jet as a model of Saturn's north-polar hexagon*. The Astrophysical Journal Letters, 806, L18. https://doi.org/10.1088/2041-8205/806/1/L18.
- Rostami, M., Zeitlin, V., & Reznik, G. (2017). *A shallow-water model for Saturn's North Polar Hexagon: jet stability and wave selection*. Icarus.
- Bach, B., Linnartz, D., Vested, M. H., & Bohr, T. (2014). *From axisymmetric flow to rotating polygons: laboratory and theory*. Journal of Fluid Mechanics.
- Bergmann, R., Tophøj, L., Homan, T. A. M., Hersen, P., & Bohr, T. (2011). *Polygon formation and surface flow on a rotating fluid surface*. Journal of Fluid Mechanics.
- Jansson, T. R. N., Haspang, M. P., Jensen, K. H., Hersen, P., & Bohr, T. (2006). *Polygons on a rotating fluid surface*. Physical Review Letters, 96, 174502.

- Tophøj, L., Mougel, J., Bohr, T., & Fabre, D. (2013). *Rotating polygon instability of a swirling free surface flow*. Physical Review Letters, 110, 194502.
- Vatistas, G. H. (1990). *A note on liquid vortex sloshing and polygonal core formation*. Journal of Fluid Mechanics.
- Marcus, P. S., & Lee, C. (1998). *Vortex core waves and polygonal patterns in rotating flows*. Physics of Fluids.
- Bastin, M. E., & Read, P. L. (1998). *Dynamics of baroclinic jet instability and equilibrated wave patterns in a rotating annulus*. Journal of Fluid Mechanics.
- Früh, W.-G., & Read, P. L. (1999). *Experiments on transition and equilibration of baroclinic waves in a differentially heated rotating annulus*. Journal of Fluid Mechanics.
- Sommeria, J., Meyers, S. D., & Swinney, H. L. (1989). *Laboratory demonstration of multiple zonal jets and equilibrated meanders from barotropic instability*. Nature, 337, 58–61.
- Read, P. L., et al. (2022). *Energy exchanges in Saturn's polar regions from Cassini observations and models* (contextual review with lab links).
- Science Magazine News (2010). *Saturn's strange hexagon recreated in the lab* (popular summary of the rotating-tank result).

16.3 Shallow-water, Rossby wave, and barotropic instability references

- Vallis, G. K. (2017). *Atmospheric and Oceanic Fluid Dynamics* (2nd ed.). Cambridge University Press.
- Pedlosky, J. (1987). *Geophysical Fluid Dynamics* (2nd ed.). Springer.
- Holton, J. R., & Hakim, G. J. (2012). *An Introduction to Dynamic Meteorology* (5th ed.). Academic Press.
- Gill, A. E. (1982). *Atmosphere–Ocean Dynamics*. Academic Press.
- Rossby, C.-G. (1939). *Relation between variations in the intensity of the zonal circulation of the atmosphere and the displacements of the semi-permanent centers of action*. Journal of Marine Research.
- Kuo, H. L. (1949). *Dynamic instability of two-dimensional nondivergent flow in a barotropic atmosphere*. Journal of Meteorology.
- Charney, J. G., & Stern, M. E. (1962). *On the stability of internal baroclinic jets in a rotating atmosphere*. Journal of the Atmospheric Sciences.

- Fjørtoft, R. (1950). *Application of integral theorems in deriving criteria of stability for laminar flows and for the baroclinic circular vortex*. Geofysiske Publikasjoner, 17(6).
- Drazin, P. G., & Reid, W. H. (2004). *Hydrodynamic Stability* (2nd ed.). Cambridge University Press.
- Yadav, R. K., & Bloxham, J. (2020). *Deep rotating convection generates the polar hexagon on Saturn*. Proceedings of the National Academy of Sciences.

16.4 DIY build references: motors, slip rings, safety, measurement tools

- Occupational Safety and Health Administration (OSHA). *29 CFR 1910.212 — General requirements for all machines (machine guarding)*.
- National Institute for Occupational Safety and Health (NIOSH). *Preventing electrocutions by undetected feedback electrical energy* (NIOSH Publication No. 88-104).
- CPWR / NIOSH. *Electrical Safety: Power (toolbox talk / training handout)*.
- International Electrotechnical Commission (IEC). *IEC 61140: Protection against electric shock — common aspects for installation and equipment*.
- IET Labs. *The Stroboscope Handbook* (strobe methods for speed measurement and aliasing pitfalls).
- Thielicke, W., & Stamhuis, E. J. (2014; updated software documentation). *PIVlab: time-resolved digital particle image velocimetry tool for MATLAB* (software + method notes).
- OpenPIV contributors. *OpenPIV / OpenPIV-MATLAB* (open-source PIV workflows and validation notes).
- Pololu Robotics & Electronics. *DC motor driver and PWM control documentation* (practical wiring, current limits, and PWM behavior).
- Adafruit Learning System. *PWM / DC motor control guides* (practical electronics patterns for hobby-scale builds).

The author has published over 4,800 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.